

AR-010-587

19981112 035

Psychological Approaches To Data
Visualisation

Michael D. Lee and Douglas Vickers

DSTO-RR-0135

APPROVED FOR PUBLIC RELEASE

© Commonwealth of Australia

DEPARTMENT OF DEFENCE
DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION

PSYCHOLOGICAL APPROACHES TO DATA VISUALISATION

Michael D. Lee

Communications Division

and

Douglas Vickers

Department of Psychology,
University of Adelaide

Electronics and Surveillance Research Laboratory

DSTO-RR-0135

ABSTRACT

The aim of 'data visualisation' is to display a body of information in a way which allows accurate and effortless human comprehension and analysis. Accordingly, the development of data visualisation techniques should be constrained by an understanding of both human perception and cognition. This report develops and examines a psychological framework for the development of data visualisation techniques based on the notion of similarity structure modelling. Through a series of case studies, a range of established approaches to data visualisation is reviewed and evaluated within this framework, and a number of suggestions for the development of new techniques is made.

APPROVED FOR PUBLIC RELEASE

DEPARTMENT OF DEFENCE



DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION

Published by

DSTO Electronics and Surveillance Research Laboratory

PO Box 1500

Salisbury, South Australia, Australia 5108

Telephone: (08) 8259 5555

Facsimile: (08) 8259 6567

© Commonwealth of Australia 1998

AR No. AR-010-587

July, 1998

APPROVED FOR PUBLIC RELEASE

Psychological Approaches To Data Visualisation

EXECUTIVE SUMMARY

In many defence-related contexts, human analysts must make decisions based on their understanding of large volumes of 'raw' information. For example, when intelligence analysts compile reports for dissemination, they must be able to search through and integrate repositories of information which have arrived from a variety of different sources. The performance of analysts on these sorts of tasks is strongly influenced by the quality of their understanding of the raw data. A broad and accurate understanding allows unusual events to be detected, changes or trends to be identified, and selective focus to be placed on those sub-sections of the information which are the most relevant for detailed analysis. Without such a high-level understanding, however, it is difficult to extract key pieces of information efficiently, or to integrate disparate collections of information into a coherent whole. Clearly, the effectiveness of the techniques which analysts use to develop their understanding of available information has an important influence upon the quality of their analyses.

Perhaps the most popular and powerful approach for examining large volumes of disparate information comes in the form of techniques generically described as 'data visualisation' techniques. These techniques aim to present data to a human using graphical displays, in ways which both accurately communicate information, and require minimal effort for comprehension. This report argues that the development of data visualisation techniques should be guided by an understanding of human perception and cognition. In this way, data visualisation can become an effective interface between a body of raw information and a human analyst. Indeed, the ultimate goal of data visualisation should be to establish a 'communication channel' between data held in artificial systems, and the humans who must deal with this information.

This report presents and demonstrates a psychological framework for developing data visualisation techniques. The framework is based on two fundamental assumptions:

- information is more meaningful to a human if it is represented in a way which is compatible with human mental representation, and
- information is more accurately conveyed to a human if it is presented in a way which is compatible with human visual processes.

The first of these observations can only be satisfied by developing an understanding of human cognition, while the second requires an understanding of human perception.

Against this psychological background, a model of human mental representation is described which allows almost any set of raw data to be converted into a representation compatible with human cognition. A variety of pre-processing methods for creating these sorts of mental models are detailed, and various types of representations which may be developed from them are described. In particular, a series of four different, but often complementary, approaches – covering spatial, featural, structural and transformational representations – are examined, and practical techniques for their generation are outlined.

The application of these different representational approaches is then examined through a series of nine case studies. The data sets involved in these studies range significantly in nature. For example, the case studies explore visualising the meaning of text documents, the patterns of confusion of Morse code signals, the voting patterns of politicians, and the perception of relationships between nations. Particular emphasis is placed on the way in which the psychological representations of the data sets may be presented, so that they are able to be immediately and accurately perceived by an observer. A number of display techniques are canvassed in these case studies, and each is subjected to critical appraisal.

Finally, on the basis of the established psychological framework, a number of suggestions are made for developing improved data visualisation techniques. Particular attention is given to the need to develop techniques for generating new types of representations, with some discussion of the ways in which these representations might appropriately be displayed. These suggestions arise from observed shortcomings in the various displays and representations presented in the case studies, and provide an agenda for future research.

Authors

Michael D. Lee

Communications Division

Michael Lee completed a Bachelor of Science in the Faculty of Mathematical Sciences between 1988–1990, Honours in the Department of Psychology in 1992, and a Ph.D. jointly supervised within the Departments of Psychology and Electrical & Electronic Engineering between 1994–1996, all at the University of Adelaide. During the final year of Ph.D. candidature, he was a postgraduate member of the Cooperative Research Centre for Sensor Signal & Information Processing. In January 1997 he commenced employment with what is now known as the Intelligent Networks Group, in the Secure Communications Branch of the Communications Division at the DSTO.

His research interests are in the areas of cognitive modelling and artificial intelligence, particularly in relation to models of human mental representation.

Douglas Vickers

Department of Psychology, University of Adelaide

Douglas Vickers completed an M.A. (Honours, Philosophy) at Edinburgh University in 1961, a B.A. (Honours, Natural Sciences Tripos) at Cambridge University in 1963, and a Ph.D. (Experimental Psychology) at Cambridge University in 1967. He took a Lectureship in Psychology at Adelaide University in 1967, and was appointed to a Readership (now Associate Professorship) there in 1979. In 1980 he spent a sabbatical year as a French Government Scientific and Professional Scholar at the Institut de Neurophysiologie, Centre National de la Recherche Scientifique at Marseilles. In 1994 he was awarded the degree of Sc.D. by Cambridge University.

His research interests are in the areas of visual perception, language, and cognition and in modelling judgment and decision processes.

THIS PAGE IS INTENTIONALLY BLANK

Contents

1	The Psychological Basis Of Data Visualisation	1
2	Cognitive Representation	3
2.1	Similarity Structure Modelling	4
2.2	Generating Similarity Measures	5
3	Modelling Similarity Structures	6
3.1	Spatial Representations	7
3.2	Featural Representations	9
3.3	Structural Representations	10
3.4	Transformational Representations	12
4	Visualisation Case Studies	13
4.1	Iris Measurements	13
4.2	Congressional Voting	19
4.3	Morse Code	20
4.4	Document Semantics	21
4.5	Zoo Animals	22
4.6	English Letters	28
4.7	Arabic Numerals	30
4.8	Drug Use	34
4.9	Nations	36
4.10	Summary	38
5	Future Prospects	38
5.1	High-Dimensional Spatial Representations	39
5.2	Integrated Similarity Structures	41
5.3	Conclusion	43
	References	44

THIS PAGE IS INTENTIONALLY BLANK

1 The Psychological Basis Of Data Visualisation

'Data visualisation' techniques aim to present data to a human in a way which both accurately communicates information, and requires minimal effort for comprehension. For this reason, graphical depictions employed in data visualisation should be constrained by an understanding of the human visual system. Any information presented as a visual display is subject to the biases and idiosyncracies of perceptual processing, and it is only by understanding these processes that visualisation techniques can be developed which do not distort the information intended to be conveyed.

Data visualisation techniques also aim to make possible human manipulation and analysis of a body of information. Accordingly, the structure of the information conveyed needs to be compatible with the representational requirements and preferences of human cognitive processes. This is one reason why data modelling techniques employed in data visualisation should also be constrained by an understanding of human memory and cognitive representation.

A second reason is that the strong interaction and inter-dependence between perception and cognition suggests that visual perception is sensitive to the structural organisation of human memory. Theories such as 'Psychophysical Complementarity' [1, 2, 3, 4] and 'Psychological Essentialism' [5, 6] argue that "selective pressures of biological evolution ... have shaped, in higher organisms, a perceptual mechanism whereby objects are represented in a way which preserves the information most essential for survival - information about the inherent properties of objects" [2]. The implication is that the veridical visual presentation of a body of information is best achieved by data visualisation techniques which represent that domain in a way which models human cognitive representation.

For both of these reasons, a principled psychological approach to data visualisation must be concerned first with the cognitive representation of information, and then with the effect of perceptual processes upon the communication of this information. It is surprising, therefore, that the development of data visualisation techniques is not routinely approached from a cognitive modelling perspective. While it is common [7, 8, 9, 10, 11] to find the need for "better integration with perceptual psychology" [11] acknowledged, the need to structure information according to the representational dictates of human memory, with some notable exceptions [12, 13], is rarely advocated.

This means, unfortunately, that many data visualisation techniques perform little or no manipulation of a body of data before an attempt at its graphical depiction is made. Sometimes this absence of pre-processing is defended on the basis of 'letting the data speak for itself', rather than imposing structure through the application of data modelling. However, this rationale, although well-intentioned, is fundamentally flawed. Humans are compulsively active and idiosyncratic constructors and interpreters of the information they receive. "We are not simply passive receptors, we actively organize and make sense of the world, and when we do so we are at the mercy of the wiring of our eyes and brains" [13]. If a data visualisation technique does not perform some form of cognitive modelling, then the human cognitive system will undertake the task of restructuring the presented information itself. In either case, the information undergoes substantial reorganisation. Put simply, if one wishes to 'let the data speak for itself', the last thing one should do is show it to a human.




Figure 1: The perceptual and cognitive basis of data visualisation. By ensuring compatibility between representation in the artificial system and human cognition, and between visualisation in the artificial system and human perception, an effective means of conveying information between the two systems may be established.

To the extent, however, that cognitive data modelling and perceptually compatible displays are provided, a data visualisation technique comes to approximate a seamless and transparent interface between a body of information and the human cognitive system. In this sense, the development of formal quantitative data models, derived from more general theories of human perception and cognition, allows for the specification of a representational 'protocol' by which artificial information processing systems and the human cognitive system may be coupled. As shown in Figure 1, the ultimate goal of data visualisation is to establish a channel between knowledge representation in an artificial system and human cognitive representation, as mediated by the perception of structured data displays.

Clearly, the integration of humans and artificial systems in this way offers enormous promise in terms of the comprehension, manipulation, and analysis of large volumes of data. The information processing capabilities of cognitive processes and artificially intelligent mechanisms are largely complementary, and their integration and interaction may well have the potential to provide capabilities which are currently exhibited by neither.

2 Cognitive Representation

The prevalent psychological view [14, 15, 16, 17, 18] is that human cognitive representation takes the form of a mental ‘conceptual’ structure, containing a large number of concepts which correspond to categories, or natural kinds, of objects in the world. This conceptual structure provides an organisational basis for human memory which fulfills at least two adaptive functions. First, the specification of a relatively small number of concepts serves to make manageable the enormously vast and complicated stream of information continually available to humans. Secondly, the concepts themselves create or preserve meaning by encompassing a group of objects which, in some sense, are capable of unitary and coherent description. In this way, the formation of human conceptual structure amounts to the formation of an adaptive and efficient modelling of the structure of the world. These two guiding principles have been aptly described as, respectively, the promotion of “cognitive economy” [16], and the maintenance of “reflections of the environment in memory” [19].

Not surprisingly, at an epistemological level, the nature of conceptual structures remains an unresolved issue subject to ongoing philosophical investigation and debate [20, 21, 22, 23, 24]. A well established and useful framework, however, considers conceptual structures in terms of their ‘horizontal’ and ‘vertical’ dimensions [16]. The horizontal dimension of a conceptual structure refers to the internal structure of mental concepts in that structure. It addresses the mechanisms by which concepts become coherent and unitary entities, and the way in which information is partitioned by the creation of these concepts. There are a number of competing formal approaches for modelling the internal structure of concepts, including the set-theoretic [25, 26], prototype [27], exemplar [28, 29] and explanation-based [30, 18] approaches, which have been described and critically compared in a relatively recent review [15].

The vertical dimension of a conceptual structure refers to the relationship between different concepts within a structure. It addresses the mechanisms by which an architectural taxonomy of concepts is created and maintained, serving to interconnect concepts which exist at different levels of abstraction. More abstract concepts are those which are more general, in the sense that they encompass a larger number of stimuli, while less abstract concepts are more selective and focussed. An important psychological notion in this regard is that of the ‘basic’ conceptual level [31], which may be conceived as ‘natural’ or ‘optimal’ balance between the efficiencies of general abstract concepts and the inherent meaning of specific concepts [32]. For example, of the concepts labelled by the words ‘furniture’, ‘chair’ and ‘bar-stool’, it is that indicated by the word ‘chair’ which is considered to exist at the basic level, while ‘furniture’ is regarded as a ‘super-ordinate’ concept, and ‘bar-stool’ is regarded as a ‘sub-ordinate’ concept. It is the meaningful interconnection of super-ordinate, sub-ordinate and basic level concepts, on the basis of their generality and common stimuli, which results in a mental conceptual structure¹. An attempt at the schematic depiction of such a structure is made in Figure 2, where concepts of different levels of abstraction are shown as spheres of different sizes, and the various conceptual interconnections across levels of abstraction are shown by connecting lines.

¹It is interesting to note that the notions of horizontal and vertical dimensions of a conceptual structure correspond almost precisely to the notions of ‘is-a’ and ‘part-of’ hierarchies in complex systems, as developed in object-oriented approaches to data modelling [33].




Figure 2: A schematic depiction of a conceptual structure.

2.1 Similarity Structure Modelling

In terms of quantitative psychological modelling, the dominant approach has been to build representations of conceptual structures using the similarity relationships existing between stimuli in a domain of interest. As early as the 1950s, a large and influential body of research proceeded on the basis that “similarity was *the* central problem of psychology” [34], and a recent overview maintains that “similarity plays an indispensable foundational role in theories of cognition” [35].

The emphasis on psychological similarity in modelling cognitive representation may be seen to stem from its relationship to the cognitive process of generalisation. In one sense, psychological similarity is simply an operational measure of the strength or likelihood of generalisation, which may be defined as the (not necessarily overt) act of treating two stimuli as if they were the same, despite the ability to discriminate between them. Thus, the process of generalisation may suggest that a red berry is poisonous because of a previous unpleasant gustatory experience with a red berry of a discernibly different hue. In this way, generalisation allows information learned in the past to be brought to bear on present concerns. Generalisation also provides the mechanism for conceptual coherence, in the sense that objects which, like the berries, are adaptively treated as the same are precisely those which should belong to the same mental concept.

Indeed, it has been argued compellingly [4] that generalisation is the most fundamental cognitive process, since all psychological processes must be considered in the context of an understanding of their operation under altered conditions, and this understanding must, itself, ultimately be founded upon an understanding of the process of generalisation. To the extent that generalisation is the fundamental psychological phenomenon, it seems likely that cognitive representation is structured in a way which reflects and facilitates

the operation of this process. Accordingly, it is reasonable to assert that measures of the psychological similarity of a set of stimuli contain much of the information needed to model a human conceptual structure of those stimuli. Certainly, such an approach is implicit in some established data visualisation approaches [36], and is consonant with the general constraint on graphical depiction identified by Kosslyn: "Use a graph only if the point is to illustrate relations" [13].

2.2 Generating Similarity Measures

In practice, similarity measures within a given stimulus domain are typically generated by considering the similarity of all possible pairs of stimuli, and are usually normalised to lie between 0 and 1. Depending upon the form and scale of the 'raw' data which constitutes the body of information under investigation, there are a variety of established means for generating indices of pairwise similarity.

Lists Of Properties

Most commonly, raw data takes the form of a series of values across m properties for a set of stimuli, which may be denoted by $\mathbf{x}_i = (x_{i1}, \dots, x_{im})$ for the i th stimulus. In this case, any monotonically decreasing function, $f_{mon}(\cdot)$, of a weighted Minkowskian distance between \mathbf{x}_i and \mathbf{x}_j provides a measure of similarity between the i th and j th stimuli. Accordingly, the similarity measure, s_{ij} , is given by:

$$s_{ij} = f_{mon} \left(\left[\sum_{k=1}^m w_k |x_{ik} - x_{jk}|^r \right]^{\frac{1}{r}} \right), \quad (1)$$

where w_k is the weight of the k th property, and r is a parameter which determines which of the family of Minkowskian distance metrics is employed. Other distance metrics, not subsumed by this general family, such as the Mahalanobis, Canberra, Bray-Curtis and Bhattacharyya distances have also been employed to derive similarity measures in this way [37].

Of course, to use metrics of this type when some of the properties have nominal or ordinal levels of data scaling, it is necessary to define uni-dimensional 'distances' between each possible pair of values for these properties. For example, if a property details an individual's nationality, it is necessary to provide some contextually meaningful quantitative measure of the difference between each possible pair of nationalities. Although it is usually possible to provide sensible definitions of this type, they typically are highly domain and property dependent.

When all of the properties have a ratio level of scaling, however, an alternative means of deriving indices of similarity is through measures of angular separation:

$$s_{ij} = f_{mon} \left(\frac{\sum_{k=1}^m x_{ik} x_{jk}}{\left[\sum_{k=1}^m x_{ik}^2 \sum_{k=1}^m x_{jk}^2 \right]^{\frac{1}{2}}} \right), \quad (2)$$

or correlation:

$$s_{ij} = \frac{\sum_{k=1}^m (x_{ik} - \bar{x}_i)(x_{jk} - \bar{x}_j)}{\left[\sum_{k=1}^m (x_{ik} - \bar{x}_i)^2 \sum_{k=1}^m (x_{jk} - \bar{x}_j)^2 \right]^{\frac{1}{2}}}. \quad (3)$$

Finally, in the special case that each of the properties x_{ik} is a binary variable, a large number of methods for generating similarity measures have been proposed, and are detailed in [37].

Measures Of Association

Probability values represent a second general form of ‘raw’ data amenable to conversion into pairwise similarity measures. Such probabilities may describe the likelihood of confusion of pairs of stimuli, their probability of co-occurrence, their frequency of substitution, or any of a range of plausible indicators of general association.

Given a matrix of pairwise probabilities, denoted p_{ij} , arising from any of these approaches, it is often convenient simply to treat each as a measure of psychological similarity. A more psychologically principled [38] approach to generating these measures is given by the relation:

$$s_{ij} = \sqrt{\frac{p_{ij}p_{ji}}{p_{ii}p_{jj}}}, \quad (4)$$

while a third option, which may have some practical advantages in terms of robustness [39], takes the form:

$$s_{ij} = \frac{p_{ij} + p_{ji}}{p_{ii} + p_{jj}}. \quad (5)$$

3 Modelling Similarity Structures

The goal of similarity structure modelling is to transform a ‘first-order’ matrix $\mathbf{S} = [s_{ij}]$, which details only how generalization operates locally between objects, into a coherent global account of the way in which generalization dictates the conceptual structure of the entire domain. In effect, similarity structure modelling imposes a theoretical framework upon a collection of ‘raw’ similarity measures in an attempt to reveal an ‘implicit’ or ‘latent’ underlying conceptual structure. Accordingly, broadly different approaches to modelling similarity structures may be defined in terms of the theoretical representational frameworks, and associated models of psychological similarity, they adopt. Following [35], at least four such approaches may be identified: spatial, featural, structural, and transformational.

3.1 Spatial Representations

The spatial approach to cognitive modelling represents stimuli as points in a multidimensional coordinate space. Therefore, if the space is of (pre-determined) dimensionality m , each stimulus is associated with a vector of the form $\mathbf{p}_i = (p_{i1}, \dots, p_{im})$. Using these vectors, the similarity between the i th and j th stimuli is modelled as a monotonically decreasing function of the distance between their representative points in the space, \hat{d}_{ij} , as follows:

$$\hat{s}_{ij} = f_{mon}(\hat{d}_{ij}) \quad (6)$$

In this way, stimuli which are more similar are afforded spatial representations which are closer together, while dissimilar stimulus pairs are placed further apart. However, it is worth noting that this approach, because of its spatial nature, is bound by the metric axioms of minimality ($d_{ij} > d_{ii} = 0$ where $i \neq j$), symmetry ($d_{ij} = d_{ji}$), and the triangle inequality ($d_{ij} \leq d_{ik} + d_{jk}$). Although it has been argued that violations of these axioms observed in empirically gathered psychological similarity data can be accounted for in a cognitive modelling context [40], it remains true that, in practice, it is often helpful to ensure that the minimality and symmetry axioms, in particular, are satisfied. Minimality may be achieved by simply ensuring that each (distinct) stimulus has a maximal self-similarity of 1, and asymmetric similarity matrices are frequently [41, 42] rendered symmetric by averaging each pair of transpose elements in \mathbf{S} .

It is also common to restrict the distance metrics of interest to the family of un-weighted² Minkowskian r -metrics³, given by

$$\hat{d}_{ij} = \left[\sum_{k=1}^m |p_{ik} - p_{jk}|^r \right]^{\frac{1}{r}}, \quad (7)$$

with a particular emphasis having been placed on the $r = 1$ (City-Block) and $r = 2$ (Euclidean) cases because of their relationship, respectively, to so-called ‘separable’ and ‘integral’ stimulus domains [45, 46]. Integral stimuli are those, such as colours, which are relatively unanalyzable, in the sense that they are not readily perceived in terms of their component dimensions. Separable stimuli, in contrast, are those in which a number of component dimensions can be considered independently, such as a set of geometric stimuli varying in size and shape. Empirically, integral stimulus dimensions may be identified by testing for filtering interference, whereby performance in attending to one dimension is affected by the other, and redundancy gains, whereby performance on one dimension is facilitated by redundant information on the other. For both separable and integral stimulus domains, Shepard and others [47, 4, 46, 48, 49] have presented compelling theoretical and empirical evidence that the monotonically decreasing function which relates similarity to distance is invariant, and is closely approximated by an exponential decay function.

²Weighted metrics are used in individual differences applications of multidimensional scaling, but are not appropriate here, since the goal is to produce a spatial representation of a *single* similarity structure.

³Although multidimensional scaling techniques have been developed [43, 44] which operate in spaces not accommodated by the Minkowskian family of metrics.

More generally, it has been argued [46] that the distinction between separable and integral stimuli may represent endpoints of a continuum rather than a dichotomy. In particular, some stimulus dimensions satisfy the filtering interference but not the redundancy gains criterion for integrality, and are sometimes termed ‘configural’ dimensions. It seems likely that domains containing stimuli of this type may be modelled appropriately using Minkowski r -metrics with an r value between 1 and 2. Although values of r greater than 2 are sometimes considered [50, 46], it is difficult to provide a psychological interpretation, in terms of component structure, for stimuli modelled in this way. Pure integrality at $r = 2$ would seem to constitute a psychological upper limit on the degree to which underlying stimulus dimensions may be combined. In contrast, the adoption of metrics with $r < 1$ has been given a psychological justification [51, 47, 46] in terms of modelling stimuli with component dimensions which ‘compete’ for attention. It seems reasonable, therefore, to conclude that there is some psychological impetus for restricting the family of Minkowski r -metrics considered in cognitive modelling to the range $0 < r \leq 2$.

Interestingly, some computational corroboration of this assertion may be found in evidence that, for every metric with $r > 2$, there is another ‘quasi-equivalent’ metric with $r < 2$ which is capable of accommodating a spatial representation with essentially the same level of error [52, 53]. In this sense, the introduction of metrics with $r > 2$ does not afford representational possibilities not available using $0 < r \leq 2$. It is, however, necessary to restrict attention further to the interval $1 \leq r \leq 2$ to preserve the metric structure of a multidimensional scaling representation, in the sense of satisfying the metric axioms.

In practice, the generation of spatial representations from a given matrix of similarity measures is most often achieved using the family of statistical techniques which may be described by the generic label ‘multidimensional scaling’ [54, 37, 55, 50, 56, 57, 58]. Most multidimensional scaling techniques employ gradient-descent or other optimisation principles to derive locations for each stimulus in a space of given dimensionality. Recent developments [59] in measuring the complexity of these representations allow for an appropriate dimensionality to be selected on the basis of the precision of the similarity data, and also provide some guidance in the selection of a distance metric.

Recently, multidimensional scaling techniques have been extended and refined by a series of neural network models such as SAMANN [60], Curvilinear Component Analysis [61], Neuroscale [62], and others [63]. The main contribution of these techniques is to introduce the possibility of learning domain specific non-linear transformations which map ‘raw’ descriptions of stimuli into appropriate positions within the derived spatial representation⁴. The Neuroscale approach has the added ability of being able to form representations which balance structural preservation, through satisfying pairwise similarity relationships, with information regarding the categorical membership of the stimuli, as supplied by a human user. Curvilinear Component Analysis also explicitly incorporates an analytical dependence upon a human user, with the suggestion being made that “this human control gives ‘more revealing’ results than ones obtained by automatic methods [in the sense that] this kind of representation helps to understand the structure of the

⁴The neural network formalism employed by these models, while sometimes convenient, is hardly necessary, and amounts to what is sometimes derided as “statistics for amateurs” [64]. This is an unfortunate state of affairs, since the connectionist philosophy [65] for which the neural network formalism was developed affords representational possibilities [66, 67] which might significantly improve data visualisation techniques.

data set and therefore to select appropriate techniques for further automatic processing” [61].

Another set of neural network techniques sometimes applied to data visualisation include the ‘Self-Organizing Map’ [68, 69, 70, 71], the ‘Elastic Net’ [72, 73], and the ‘Generative Topographic Map’ [74], all of which derive representations based upon the notion of preserving a quantized neighborhood topology. The effect of quantization, however, is to limit the resolution of the spatial representations of stimulus domains, without any obvious compensating benefit in terms of data modelling flexibility, efficiency or accuracy [75]. Indeed, in some sense, the basis for these techniques appears to be more physiological than psychological and, with the exception of the ‘Semantic Map’ [76], they have not been applied to the modelling of human conceptual structures.

3.2 Featural Representations

Under the featural approach to cognitive modelling, stimuli are represented in terms of the presence or absence of a set of discrete features or properties. If m such features are identified, each stimulus may again be associated with a vector of the form $\mathbf{f}_i = (f_{i1}, \dots, f_{im})$. The nature of the feature set, however, means that the f_{ik} values assume one of an enumerable set of values. Using this type of representation, the general ‘contrast’ model [77] assumes that the similarity between the i th and j th stimuli takes the form

$$\hat{s}_{ij} = \theta f_{mon}(\mathbf{f}_i \cap \mathbf{f}_j) - \alpha f_{mon}(\mathbf{f}_i - \mathbf{f}_j) - \beta f_{mon}(\mathbf{f}_j - \mathbf{f}_i), \quad (8)$$

where θ , α and β are positive weighting parameters, $\mathbf{f}_i \cap \mathbf{f}_j$ denotes the features common to the i th and j th stimuli, and $\mathbf{f}_i - \mathbf{f}_j$ denotes the features present in the i th, but not the j th, stimulus. In this way, the similarity of a pair of stimuli is modelled in terms of a comparison of the number of features both have in common with the number of features evident in one only.

Additive clustering models [78, 79, 80, 81, 82, 83] constitute perhaps the simplest, and certainly the most thoroughly developed, specific realisation of the general contrast model. Under this approach, it is assumed that the domain features are binary in nature, so that

$$f_{ik} = \begin{cases} 1 & \text{if object } i \text{ has feature } k \\ 0 & \text{otherwise} \end{cases}, \quad (9)$$

which allows the similarity between the i th and j th stimuli to be modelled as the sum of the weights of the features common to both, as follows:

$$\hat{s}_{ij} = \sum_k w_k f_{ik} f_{jk}, \quad (10)$$

where w_k is the saliency weight of the k th cluster.

It is generally recognized that the binary nature of the feature variables makes the derivation of additive clustering representations a difficult optimization problem and, accordingly, a wide variety of extraction techniques have been proposed, including mathematical programming [78], qualitative factor analytic [81] and probabilistic expectation-maximization [83] approaches. While all of these techniques have shortcomings, it is probably fair to suggest that they generally achieve sufficiently good minima to derive models of substantial theoretical and practical utility.

Importantly, all of these approaches view the relationship between the given set of stimuli and the derived features as being entirely unconstrained. Unlike partitional clustering approaches, which place each object in only one cluster, additive clustering allows each object to belong to any number of clusters. Furthermore, unlike hierarchical clustering, the additive clustering approach places no constraints upon the set of objects which may be encompassed.

This state of affairs is summarised by Figure 3, which shows (a) partitional, (b) hierarchical, and (c) additive clustering of a set of points, and their corresponding representational interpretation. The partitional approach effectively forms a collection of disjoint sets, or equivalence classes, of the points, and the hierarchical approach, because of its strictly nested nature, corresponds to the construction of a tree, in which the points are placed at terminal nodes. The additive clustering approach, however, amounts to the derivation of a set of binary features or properties, and the description of each of the points in terms of these features.

Although it seems likely that this flexibility is one of the keys to the success of the additive clustering approach [82], it also introduces the possibility of constructing over-parameterised and unconstrained models of cognitive representation. In this regard, recent developments [84] in measuring the relative complexity of different additive clustering models allow for the derivation of models with only the level of complexity warranted by the precision of the similarity data.

In principle, these theoretical results also allow constraints to be placed upon additive clustering models so that, where appropriate, they take the form of inherently simpler partitional or hierarchical cluster structures. The development of practical feature extraction techniques which implement such constraints should be a priority for future research, since it would allow a general method for deriving binary valued features under an additive similarity model. In the meantime, however, there are a wide variety of specialised techniques for generating both partitioned featural representations, such as k -means clustering, and various tree-based similarity structures, such as additive trees and ultra-metric trees [85, 86, 87].

3.3 Structural Representations

The structural approach to modelling mental representation is more general than either the spatial or featural approaches, and is “well suited for comparing things that are richly structured rather than just being a collection of coordinates or features” [35]. Under this approach, stimuli are represented in terms of generic memory structures, such as frames [88], scripts [89] and schemata [90, 67]. Using these types of sophisticated memory structures, complicated similarity relationships between stimuli – particularly those

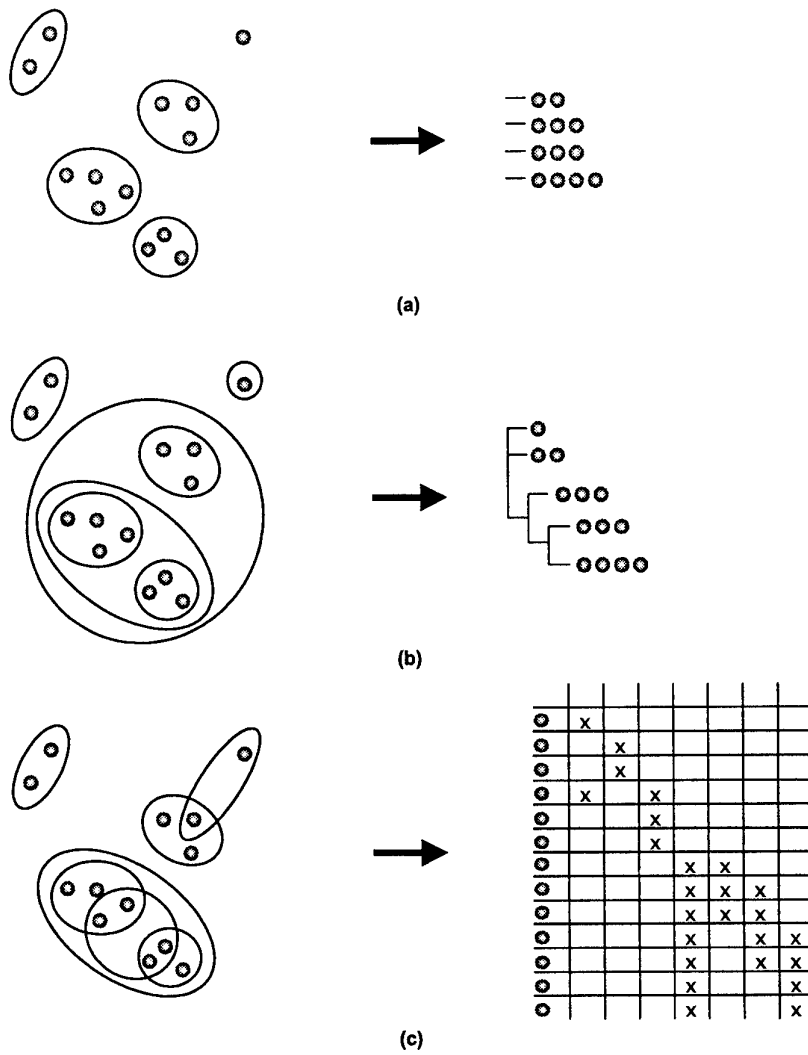


Figure 3: Typical cluster structures, and their corresponding representational interpretations for (a) partitional, (b) hierarchical, and (c) additive clustering approaches.

arising from the correspondence or matching of functional parts – may be accommodated within a formal representation. Given this capability, it is not surprising that the quantitative psychological modelling of the important cognitive process of analogical reasoning⁵ employs structural representations [93].

Unfortunately, the richness of structural models of mental representation makes their abstraction from measures of pairwise similarity, at best, a difficult problem or, more usually, renders it an under-determined problem. Accordingly, the structural approach does not have techniques such as multidimensional scaling and additive clustering for deriving conceptual structures from measures of similarity. Rather, structured domain ontologies are usually developed by alternative, sometimes unspecified, means, and these representations are simply assumed in subsequent cognitive modelling. As has previously been argued [94], this practice raises doubts that “a poor representation will often doom the model to failure, and an excessively generous representation may essentially solve the problem in advance” [95], and also limits the applicability of the structural representational approach to data visualisation. One of the primary motivations for the development of data visualisation techniques is to facilitate exploratory data analysis where, by definition, no detailed understanding of a body of information is available. Even if, in the context of domain specific analysis, some schematic representational template may be formalised, the immaturity of techniques for modelling stimuli within this framework on the basis of similarity judgments renders the general application of the structural representational approach infeasible at present.

3.4 Transformational Representations

The transformational approach to cognitive modelling is based upon the general theoretical notion that “it is not objects, but their transformations which are primary” [96]. Whereas the spatial, featural and structural approaches may be identified with a passive, object-oriented epistemology focussed on the representation of *things*, the transformational view enacts an epistemology based on the representation of active *processes*. Under this latter view, as articulated by Vickers’ proposed ‘Erlanger’ program for psychology [97, 98], mental representations correspond to the parameters of some form of reconstructive process, and conceptual structures arise from a consideration of the geometric and topological invariants of these processes. Theoretical and empirical support for the transformational approach may be found in studies of ‘apparent motion’, including particularly those involving mental rotation [99, 96, 100, 48], studies of ‘representational momentum’ [101], and various other the symmetry-based [102] and group-theoretic [103] psychological models.

In terms of similarity structure modelling, the transformational approach assumes that the psychological similarity of two stimuli is given by some monotonically decreasing function of the number of cognitive ‘operations’ required to transform one stimulus so as to be identical to the other. If each stimulus is given general vector representation of the form $\mathbf{r}_i = (r_{i1}, \dots, r_{im})$, the similarity between the i th and j th stimuli may be expressed as:

⁵Hofstadter has made a compelling case for the pivotal role played by analogical reasoning in underpinning ‘intelligent’ human behaviour [91, 92].

$$\hat{s}_{ij} = f_{mon}(\|\mathbf{r}_i, \mathbf{r}_j\|_{\mathcal{T}}), \quad (11)$$

where $\|\mathbf{r}_i, \mathbf{r}_j\|_{\mathcal{T}}$ denotes the number of representational transformations, drawn from a set \mathcal{T} , applied in converting \mathbf{r}_i into \mathbf{r}_j .

As with the structural approach, there is no established set of techniques for abstracting transformational representations from an arbitrary matrix of pairwise similarities. While the recently developed technique known as ‘Trajectory Mapping’ [104] develops transformational representations of stimulus domains, it requires information not provided in a standard matrix of pairwise similarities. In particular, Trajectory Mapping operates upon sequences of transformations of stimuli, as completed by a human. In some circumstances, it may be appropriate to conceive of ‘raw’ data in the form of pairwise probabilities as transitional probability measures. It is more difficult to conceive of reasonable ways in which such sequences could be generated from ‘raw’ data in the form of lists of properties or features. Indeed, in one sense, the input required by Trajectory Mapping directly provides - merely in an over-specified way - the transformational representation which is sought, and it is simply the removal of these redundancies which is accomplished by the technique itself.

The motivating framework of Trajectory Mapping, however, that of recovering the featural parameterizations used to traverse a non-homogenous conceptual space, provides some guidance as to how transformational representations may be abstracted from similarity structures. Using featural extraction techniques such as those associated with additive clustering, it seems natural and reasonable to associate featural manipulation with the transformations in the set \mathcal{T} . In other words, a transformational trajectory between two stimuli may be established if those stimuli differ by, say, one feature only. More generally, a host of graph theoretic measures can be applied to the featural characterisation of a similarity structure, since the discrete form of this representation permits the definition of various types of similarity-based connectivity between stimuli.

4 Visualisation Case Studies

4.1 Iris Measurements

The first case study involves the (in)famous ‘Iris’ database [105], which details the sepal length, sepal width, petal length and petal width of a total of 150 irises. 50 of the irises are identified as ‘iris-setosa’, 50 as ‘iris-versicolor’, and 50 as ‘iris-virginica’. This stimulus domain provides an example of ‘raw’ data in the form of a quantitative list of properties. In this case, each iris is defined by a four-dimensional, continuous valued vector of the height and width measurements made of its sepal and petal.

‘Raw’ Glyph Visualisation

It is possible to provide a graphical depiction of each of these four-dimensional vectors of ‘raw’ data, using the family of display methods generically known as ‘glyphs’ [10]. In




Figure 4: Various glyph constructs: (a) 'Fortson' glyph, in which values are represented by the shading of the bars, (b) 'Chernoff' face, in which values are represented by caricatured facial characteristics, (c) 'Star' glyph, in which values are represented by the length of rays, (d) 'Profile' glyph, in which values are represented by a plotted data series.

essence, glyphs provide a means of displaying items of multivariate data by defining an association between scalar components of an item and various perceptual features. For any given item, a corresponding glyph is constructed simply by displaying the perceptual amalgam of these features in accordance with the data of that item.

An enormous variety of glyph constructs have been proposed for data visualisation, four of which are reproduced in Figure 4. It seems clear that glyph construction is effectively limited only by one's imagination⁶ and, indeed, general techniques for the construction of arbitrary mappings between data properties and sets of pre-determined perceptual features have been developed [106]. With the exception of 'Chernoff faces' [107], however, glyph visualisations are rarely related to, constrained by, or evaluated against, the cognitive processes required for their apprehension. As shown in Figure 4, the Chernoff face approach depicts items as caricatured faces, using the values of different item properties to control such perceptual features as the width between the eyes, the length of the nose, and so on. A cognitive advantage is claimed for this form of depiction on the basis of heightened human sensitivity to perceptual differences in facial structure and expression, presumably as a result of evolutionary adaptation. The obvious difficulty with this line of reasoning is that any benefit in 'cognitive resolution' attained in this manner is counteracted by an entrenched and unconscious contamination of the affective aspect of comprehending facial depictions. As found empirically [108], Chernoff faces come with a plethora of emotional and social baggage which is difficult to quantify, but has the ability to affect the way in which information they are intended to convey is actually perceived and conceived.

A range of affectively neutral glyph constructs, such as the 'fortson', 'star' and 'profile' glyphs shown in Figure 4, have been proposed, and maintain some interest in the context of data visualisation. It is difficult, if not impossible, to provide a definitive evaluation of the relative merits of different glyph approaches, given the lack of constraints upon their construction and the absence of rigorous empirical comparison. However, the observation that lengths and directions of short line segments are perceived somewhat more accurately than gray scale colouring [109] tends to favour 'star-like' approaches over 'fortson-like' approaches. It also seems clear that the use of colour rather than a gray scale is fraught with

⁶or graphics libraries, whichever is the smaller.




Figure 5: Nine examples of 'loopy' glyphs, arranged on a 3×3 grid.

peril [110], if not because of our incomplete theoretical understanding of the complicated and subtle way in which colour is perceived [111, 48], or because of the significant conflict and contradiction evident in empirical evaluations [13, 112], then certainly because of the strong likelihood of important and uncontrollable individual differences across humans.

While some suggestions have been made regarding the appropriate designation of data components to the rays of star glyphs [109], usually on the basis of expected or observed correlations between data components, these approaches are generally relatively immature. This is unfortunate, since many of the profile-like glyph constructs, including the 'parallel-coordinates' approach to displaying multivariate data, have significant dependency upon the ordering of data components. Indeed, this dependency is so severe as to transcend issues of perceptual or cognitive effect, since the ordering dictates the actual presence or absence of physical markings.

Of course, these markings could be used to advantage, as demonstrated by the 'loopy' glyphs [113] shown in Figure 5. These 9 glyphs, arranged on a 3×3 grid, are generated by enclosing a subset of an underlying 4×4 grid of circles. As is evident from Figure 5, a wide variety of different perceptual forms are created on the basis of which circles are included in the subset. A natural application of these glyphs arises in relation to the

featural representational approach, by which sets of binary features could be associated with each of the circles. In particular, if criteria existed for assigning related sets of features to local neighbourhoods of circles, their co-occurrence would result in perceptually simple ‘blocks’ being evident in the glyphs. In this way, the perceptual organisation of displays such as Figure 5 would reflect higher-order conceptual groupings evident in a stimulus domain. Clearly, this is a desirable state of affairs, and highlights the need to develop more sophisticated methods for ordering the components of glyph displays.

Meanwhile, however, it is probably prudent to adopt the approach of star-like glyphs, in which the relative perceptual proximity of all rays places some limits on the effects of arbitrary assignment. In any case, star glyphs have been favourably compared to profile glyphs because of their visual ‘impact’ [109], and some summative empirical comparisons have also found star glyphs to be superior to more conventional graphing techniques, such as bar graphs [114].

Despite these advantages over fortson-like and profile-like constructs, the star-like approach retains several shortcomings. Findings suggesting the inaccuracy of human perception of area, as formulated in terms of Stevens’ classic psychophysical power law [115], for example, impose what have been described as “moderately severe limitations” [116] upon the ability of a star glyph to convey quantitatively precise information. In addition, the perceptual judgment of the length of a ray varies according to the orientation of that ray [117], with horizontal and vertical rays, in particular, being more accurately perceived than oblique rays. The first of these deficiencies is readily countered simply by removing the lines which form enclosed areas by connecting the rays. This modification reduces the star glyph to what is sometimes termed a ‘ray’ glyph, and provides a useful representative example of the general glyph approach.

To this end, Figure 6 presents the raw iris data using the ray glyph scheme. Each glyph on the 15×10 grid correspond to one iris, and, within each glyph, each iris measurement is represented by the length of a ray. This visualisation of the entire iris data set demonstrates several general points relating to glyph visualisations. First, contrary to the conclusion drawn by [118], it seems clear that this approach has some capability in terms of the number of stimuli and stimulus dimensions which can be displayed at any one time. Figure 6 graphically presents every measurement of every iris in the stimulus domain. Of course, the depiction of perceptually discernible quantitative information is limited by the detail required by individual glyphs, but it remains true that the glyph based approach fares reasonably well in this regard.

Unfortunately, however, it also seems clear that the glyph approach may result in the construction of displays which require considerable cognitive effort for comprehension. While Figure 6 may faithfully present the information about the various iris measurements, the partitioning of the stimulus set into the three iris types is not a trivial process, and is certainly not an ‘automatic’ or ‘unavoidable’ consequence of the visualisation. One hindrance in this regard is the lack of meaningful spatial organisation of the display. Given the widely asserted primacy of spatial proximity as a bearer of perceptual information [119, 109], the arbitrary placement of the glyphs on the grid limits the possibility of readily conveying a coherent global structure. There appears to be some considerable weight to the conclusion that, although using the glyph approach to visualisation “is sometimes an effective way to pack much information into a single display, readers typically

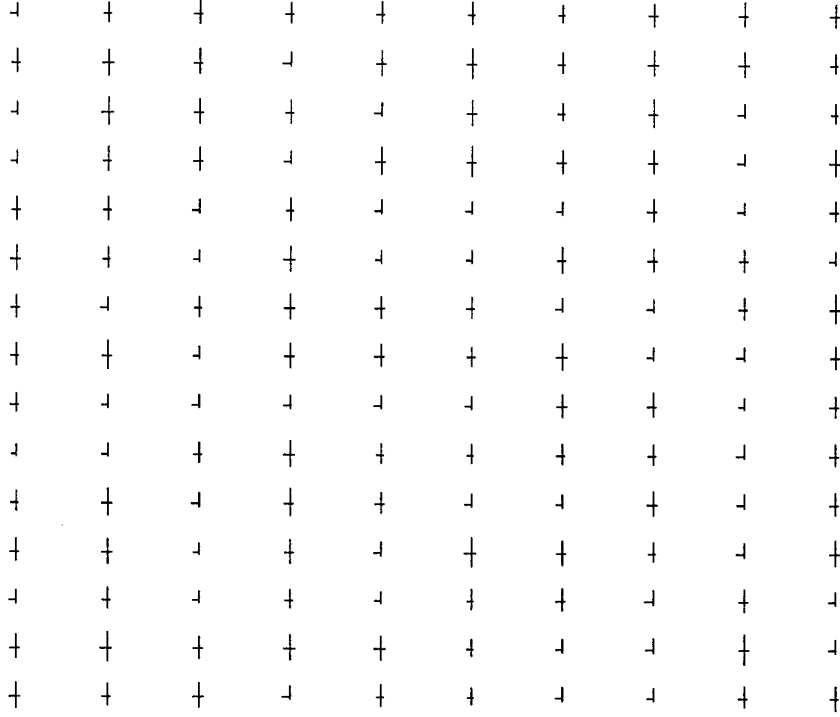


Figure 6: Ray glyph visualisation of the Iris domain.

need considerable practice before being able to decode such displays easily” [13] and, subsequently, that these techniques are “better suited for informal clustering and spotting peculiar points” [109].

Spatial Visualisation

The spatial approach to similarity structure modelling is, by its very nature, ideally suited to spatial graphical depiction. To allow the derivation of such a representation, a matrix of pairwise similarity measures was generated from the vectors of raw iris measurements, using the unweighted Euclidean version of Equation 1 and an exponential decay functional form. After subjecting this similarity matrix to multidimensional scaling, which essentially amounts to an exercise in topological preservation across spaces with different dimensionalities, a complexity analysis [59] indicated the appropriateness of a two-dimensional representation which explained 98.3% of the variance of the raw data⁷.

This two-dimensional representation is shown in Figure 7, with each iris shown by a label indicating its type. The clustered structure of the stimulus domain is readily perceived, with the iris-setosa being spatially separated, and the iris-versicolor and iris-virginica occupying essentially disjoint regions of the remaining cluster. While it is true that if the various labels were not displayed – as was the case for the glyph visualisation in Figure

⁷In general, the variance explained by a set of approximations $\{\hat{x}_i\}$ to a set of true values $\{x_i\}$ is given by $1 - \sum_i (\hat{x}_i - x_i)^2 / \sum_i (x_i - \bar{x})^2$, where \bar{x} is the average of the set of true values.




Figure 7: *Spatial visualisation of the Iris data. Iris-setosa stimuli are shown with the letter 's', iris-versicolor with 'v', and iris-virginica with 'g'.*

6 – the distinction between the iris-versicolor and iris-virginica would not be immediately perceptually obvious, this is a property of the data, rather than the visualisation. An inspection of the original physical measurements indicates that, although there is a reliable difference taken across the two classes as a whole, the division is not clear-cut at the boundary.

Figure 7 faithfully conveys this state of affairs because it is underpinned by an accurate representational characterisation of the raw information, and this representation is itself amenable to straightforward graphical depiction. It is important to note that it is the appropriateness of a two-dimensional spatial representation which facilitates the unproblematic transfer from spatial representation to spatial visualisation. Most data visualisation techniques deal with media, such as printed pages or computer screens, which are inherently two-dimensional, and it is simple to equate representational and presentational dimensions. In fact, this relationship is so natural and immediate that it is easy to forget the distinction between the abstract representation and its graphical depiction. Attempting to produce veridical presentations of spatial representations of higher dimensionality, however, is much less straightforward. The next three case studies further examine the data visualisation capabilities of the spatial approach when two dimensions prove to be sufficient, before stimulus domains which are not satisfactorily characterised in this way are tackled.




Figure 8: Visualisation of the two-dimensional spatial representation for the Congressional voting domain. Senators are labelled by their surnames, and Republican or Democratic party affiliations.

4.2 Congressional Voting

The congressional voting case study involves data detailing the voting patterns of 15 senators from New Jersey on 19 environmental bills [120]. Co-occurrence was employed as a measure of similarity simply by counting the number of times each pair of senators voted the same way across all of the bills.

Spatial Visualisation

A multidimensional scaling analysis of the similarity matrix found a two-dimensional spatial representation, explaining 89.3% of the variance of the data, to be appropriate, and is shown in Figure 8. This visualisation suggests that the voting patterns of the Democrats are more uniform, or disciplined, than those of the republicans. It also readily reveals what might be regarded as the ‘renegade’ voting behaviour of the Republican named Rinaldo. Clearly, this senator’s spatial representation asserts an alignment with what are generally Democratic voting patterns. A close examination of the actual voting details bears out these conclusions. Furthermore, both the Democrat named Thompson, and the Republican named Sandman are visually perceived as ‘outliers’ which, again with reference to the raw voting data, can be seen to be caused by their relatively greater rates of abstention from voting.




Figure 9: Morse Code numeral visualisation. The actual code of each stimulus, in terms of dots and dashes, is appended to the numeral label.

4.3 Morse Code

The third case study examines confusion data involving the identification of Morse code numbers [121], which were treated as indices of psychological similarity without modification.

Spatial Visualisation

A multidimensional scaling analysis was conducted on this stimulus domain which again revealed the appropriateness of a two-dimensional representation, explaining 90.1% of the variance in the data. In this case, however, strong evidence was found for the adoption of the City-Block distance metric in accordance with the presumably separable structure of the stimulus domain. Accordingly, it is to be expected that the associated spatial visualisation, shown in Figure 9, presents axes which are amenable to psychological interpretation.

Indeed, it is reasonably easy to infer from Figure 9 that human performance in perceiving a Morse code numeral may be explained in terms of whether it begins with a dot or dash, and the relative proportion of dots and dashes across the whole code sequence. Since the spatial proximity of the stimuli conveys their relative likelihoods of confusion, the dispersion of the numerals along the horizontal axis indicates that codes with the same proportion of dots and dashes are more often confused. Meanwhile, dispersion along the

vertical axis shows that codes started with a dot are confused with other beginning with a dot, and that the same situation holds for initial dashes.

4.4 Document Semantics

The next case study involves a textual domain containing 82 documents collected from the Internet, as previously explored in [122]. A total of 53 of these were papers from the 1995 Neural Information Processing Systems conference [123], 21 were media releases associated with the Olympic Games to be held in Sydney in the year 2000, and the remaining 8 papers were match reports involving the 1997 Australian cricket tour of England.

Each document was subject to the well-established ‘n-gram’ orthographic approach to analyzing the semantic or contextual structure of textual information [124] which, in essence creates a ‘document vector’ containing the relative frequency of all occurring sequences of n characters within the text. More specifically, a 5-gram analysis was conducted in which only 27 characters – the 26 Roman letters and the space character – were considered.

Spatial Visualisation

Following [124], pairwise indices of document similarity were generated by measuring the (interior) angle between the associated document vectors. The resultant similarity matrix was then subjected to a multidimensional scaling analysis, which provided some evidence for the appropriateness of a two-dimensional Euclidean representation⁸. The resultant spatial visualisation is shown in Figure 10, in which the documents derived from the three different sources are clearly separated into three clusters.

It is tempting to conjecture that further semantic information is available in this display, given that the relative diversity of each cluster accords well with intuition. In particular, the fact that the media releases involve a broader range of issues than the other two narrow specialist domains is reflected by the relative spatial diversity of its associated cluster⁹.

What this case study does demonstrate, together with the iris measurement study, is the capability of spatial data visualisation to display relatively large stimulus domains. In Figure 7 and Figure 10, detailed relational information concerning 150 and 82 stimuli, respectively, is conveyed without sacrificing clarity. Furthermore, using simple points to represent stimuli would allow perhaps an order of magnitude increase in the number of stimuli which could be depicted in this way. Encouragingly, it has been found that the

⁸The strict application of Bayesian complexity measures to determine dimensionality and metric structure is somewhat problematic in this case, because the precision of the similarity data is difficult to quantify. Therefore, to some extent, the ends – in the form of Figure 10 – have been used to justify the means. The spatial representation accounts for only 61.4% the variance of the data, but nevertheless seems to capture the most important components of this variation.

⁹One might even be tempted to conclude that cricket writing is one-dimensional on the basis of Figure 10, although this may be an artefact of the match reports in question largely being concerned with English victories.




Figure 10: Spatial visualisation of the document domain. Documents drawn from the academic conference are depicted as squares, documents drawn from the 2000 Olympics media releases are shown as circles, and documents drawn from the Ashes cricket reports are shown as triangles.

comprehension of spatial data displays can become more accurate as the number of representative points is increased [125], and that the time taken to achieve this understanding is, at worst, only slightly affected by increases in the number of points [126].

4.5 Zoo Animals

The next case study involves the ‘Zoo’ database [127], which was modified¹⁰ to give information regarding the presence or absence of 14 binary features (labelled ‘hair’, ‘feathers’, ‘eggs’, ‘milk’, ‘airborne’, ‘aquatic’, ‘predator’, ‘toothed’, ‘backbone’, ‘breathes’, ‘venomous’, ‘fins’, ‘tail’, and ‘domestic’) across 99 animals. Of the 42 unique featural descriptions amongst the set of 99 animals, one of each was selected at random for retention in the final stimulus set.

¹⁰The ‘animal’ labeled ‘girl’ in the original database was discarded, as was a surplus copy of the animal ‘frog’. In addition, the non-binary features ‘legs’, ‘name’ and ‘type’ were removed, as was the binary valued ‘catsize’.



Figure 11: Ray glyph visualisation of 'raw' featural representation of the zoo domain.

Featural Visualisation

Although the featural descriptions of the animals are 'raw' in the sense that they have not been explicitly derived through similarity structure modelling, they differ from the raw measurements of the iris domain in the sense that they have some psychological plausibility as mental representations. Indeed, meaningful constituent features such as 'aquatic' and 'venomous' are precisely the type of representation one would hope to derive in an additive clustering model of a zoo animal similarity structure. To examine the effect of this mental credibility upon glyph-based visualisation, Figure 11 presents a ray glyph visualisation of the zoo domain, using the same format as Figure 6 applied to the raw iris data.

As with the iris domain, the comprehension of the glyph visualisation seems effortful, but not impossible. It seems likely that, with experience, the presentation of stimuli in a structured glyph display, if underpinned by a meaningful cognitive representation, could be effective in conveying information to a human. It has been suggested [128] that periods of exposure allow the cognitive construction of 'prototypes' or 'schemata' capable of instantiating the anticipatory priming required for effortless comprehension, a notion for which there is some confirmatory empirical evidence [129, 130]. Indeed, the immaturity of guidelines for the construction of a sufficiently 'rich' perceptual combination of features is of concern in this regard. For example, it is worth noting the lack of consideration given to rotational symmetries in Figure 11. The importance of symmetry in determining perceptual, and perhaps cognitive, similarity is one with a long history [131, 45], and

an accepted currency [102, 103, 48] within psychology. That the method of ray glyph construction ignores this possibility seems likely to result in animals such as ‘seawasp’ and ‘tortoise’ which are featurally different, but happen to have glyph presentations which are almost equivalent under rotation, as being perceived and conceived as more similar than is intended or warranted.

Spatial Visualisation

To allow the derivation of a spatial representation of the zoo domain, measures of similarity between each pair of animals were again generated according to Equation 1 using an exponential decay functional form, as applied to an unweighted distance metric. It should be noted that, in this case, the binary nature of the featural properties means that the derived target distances are the same across the entire family of Minkowskian metrics. Indeed, the metric space might be more naturally characterised as a Hamming space. In any case, after applying multidimensional scaling within Euclidean spaces, a three-dimensional spatial representation, again explaining more than 89.7% of the variance of the data, was found to be appropriate. Effectively, therefore, the multidimensional scaling process involves ‘topological’ preservation across spaces differing in both their dimensionality and their metric structure.

A perspective visualisation obtained by projecting the three representational dimensions onto the two presentational dimensions is shown in Figure 12, and seems to be reasonably effective. There is, however, some cause for concern in relation to data visualisation techniques which employ this type of perspective display. While there are a number of visual cues, such as the vertical ‘reference’ or ‘drop’ lines used in Figure 12, which can be exploited to convey depth effectively on a two-dimensional medium, their relative impoverishment has resulted in the stricture “do not use three-dimensional perspective to communicate precise information” [13] being generally accepted. For this reason, there is a considerable body of research examining the various biases induced by perspective displays [132, 133].

More esoterically, a recurring theme in the work of artists such as M.C. Escher is that of the “conflict between the flat and the spatial” [134], with the basic message being that “no matter how cleverly you try to simulate three dimensions in two, you are always missing some essence of three-dimensionality” [91]. It is not easy to be definitive with regard to the extent to which this in-principle deficiency must be weighed against the practical attempts of perspective displays to convey useful and meaningful information. The observation that, “although three-dimensional graphs are rarely as confusing as Escher drawings, they sometimes lead in that direction” [13], suggests that this conflict should be recognised and minimised, rather than avoided. The conclusion that “three dimensions can be depicted well enough on a two-dimensional surface – but the process is not perfect” [13] is probably a balanced one.

Some confirmation of the relative merits of attempting to portray three dimensional spatial structures through perspective projections is gained by examining alternative approaches, in which other display features are used to present a third representational dimension. A variety of static features, such as length, angle, size, colour, texture, shape and gray-scale shading, have been suggested for this purpose. Accordingly, Figures 13, 14




Figure 12: Visualisation of three-dimensional spatial similarity structure of the zoo animals.

and 15 employ, respectively, size, orientation, and gray-scale to convey the third spatial dimensional of the zoo domain representation.

Intuitively, it seems reasonable to claim that, although Figures 13, 14 and 15 are amenable to interpretation, they require cognitive effort which is at least as great as that demanded by the perspective display in Figure 12. This appears to be particularly true of the orientational approach, which is difficult to interpret. The relatively better results given by the use of size and shading seems likely to be because they constitute natural visual cues for the conveyance of depth, fulfilling much the same role as the drop lines in the perspective display. This means, in turn, that both size and shading could be incorporated into a perspective display, in which case the perspective approach is likely to be superior to all three of the presented alternatives.

Of course, the issue of the relative merits of various perceptual features for conveying additional spatial dimensions is amenable to principled psychological investigation. Establishing base-rates for the accuracy with which various features are perceptually perceived and cognitively judged should impact upon the development of data visualisation techniques. To this end, based on a mixture of theoretical notions and empirical findings, a crude ordering of the efficacy of various features in conveying quantitative variation has




Figure 13: Visualisation of spatial structure shown in Figure 12, using two spatial dimensions and a size dimension.

been suggested [135] which gives spatial location as the most effective, followed by length, orientation, size, volume, density, colour saturation, and colour hue. Unfortunately, further examination of relevant studies in this area tends to affirm the less prescriptive conclusion that “the accuracy of simple judgments of amount is not the only factor you should consider when deciding how to display data” [13]. In particular, the representational effects of different features have been shown to be susceptible to practice effects [136], task demands [130, 137], and the perceptual features by which they are surrounded in a multi-faceted display [138]. There is, therefore, considerable scope for ongoing theoretical and empirical evaluation of data visualisation techniques in the context of exploratory data analysis, and perhaps even in terms of specific target domains of interest.

Transformational Visualisation

The featural representation on which the zoo domain is founded allows an exploration of the ‘featural connectivity’ notion arising from the transformational approach to cognitive representation. Figure 16 simply presents the best fitting two-dimensional spatial representation of the animals obtained from the previous multidimensional scaling analysis, and introduces connecting lines to join those animals which differ by only one feature. In this way, a transformational representation of the stimulus domain, which is essentially a graph – or set of connected graphs – can be presented visually.

There are a number of established techniques, reviewed in [139], which present graph-




Figure 14: Visualisation of spatial structure shown in Figure 12, using two spatial dimensions and an orientation dimension.

structures on the basis of a series of 'aesthetic principles', rather than imposing the graph upon a spatial representation. Many of these 'principles' seem to have been proposed on an *ad hoc* basis, but nonetheless often appear to be reasonable, and might even be amenable to being placed on a more psychologically satisfactory foundation through a consideration of Gestalt principles of perceptual organisation [140]. For example, the notion that connecting lines should preferably not be drawn to intersect seems related to the Gestalt principle of good continuation. However, graph layouts produced by the multi-objective optimisation of a conglomerate of these aesthetic strictures are generally far from appealing, particularly when applied to real rather than highly artificial similarity structures, and little or no evidence has been provided that they actually facilitate the communication of information. Certainly, aesthetic graph layout techniques have not been considered in terms of their cognitive representational implications. Presumably, this state of affairs arises because different aesthetic principles sometimes conflict, and their relative perceptual saliences are poorly understood, with the result that attempts at their simultaneous satisfaction often results in none, rather than some or all, of all the principles being evident in a display.

Perhaps the one instance in which aesthetic graph layout techniques are of some value is when suitable spatial representations of a domain are difficult to derive because of an impoverished similarity structure. In particular, binary measures of pairwise similarity are unlikely to be able to be satisfactorily represented in low-dimensional coordinate spaces, and induce idiosyncratic behaviour from several multidimensional scaling techniques [141].




Figure 15: Visualisation of spatial structure shown in Figure 12 using two spatial dimensions and a gray-scale dimension.

Although, in principle, it should always be possible to provide more detailed measures of pairwise similarity, some experimental methodologies, particularly in the context of social ‘networks’, tend to limit themselves to simple binary observational measures of relationships. Unless further assumptions are made about the nature of these relationships, and more detailed similarity indices are derived on this basis [54], aesthetic graph layout approaches may be superior to those employed in Figure 16 for displaying transformational representations of the domain.

4.6 English Letters

Empirical confusion probabilities for the 26 uppercase English letters, as reported in [142] form the basis of the next case study. Basically, these measures reflect confusions made in terms of the perceived visual structure of the letters. As with the Morse code stimuli, these probabilities are themselves treated as measures of psychological similarity.

Spatial Visualisation

A multidimensional scaling analysis of this similarity structure found justification, even when relatively severe assumptions regarding the precision of the data were made, for the inclusion of at least five spatial representational dimensions, accounting for 67.4%




Figure 16: Visualisation of a feature based transformational representation for the zoo domain.

of the variance in the data. Figure 17 attempts to display the best fitting five dimensional representations using two spatial dimensions, orientation, size and gray-scale shading. If any meaningful information regarding the structure of the letter domain is conveyed by this approach, then it is certainly only gleaned by the effortful application of considered (and serial) cognitive processes. Figure 17 does not immediately convey the information presumably contained within the derived spatial representation, and serves to highlight the deficiencies of this approach first observed in relation to the zoo domain.

An alternative means of displaying high-dimensional spatial representations of this type is through producing exhaustive planar projections of the coordinate space [87], to form what is sometimes¹¹ referred to as a 'draftsperson's plot'. Figure 18 presents all $5 \times 4 \div 2 = 10$ unique combinations of the five-dimensional spatial representation for the letter domain. This visualisation is, because of its inherently fragmented nature, entirely unsuited to the apprehension of a global domain structure, and does not even seem to be superior to Figure 17 in terms of the accurate communication of more local information.

However, the exhaustive planar projective approach is likely to provide a very effective means of visualising high-dimensional spatial representations which are appropriately assumed to be psychologically separable. As noted in relation to the Morse code domain, asserting the operation of separability through the application of the City-Block distance metric results in the derivation of spatial representations with interpretable axes. In this case, the planar projection approach displays the relationship between stimuli within

¹¹after gender neutering




Figure 17: Visualisation of five-dimensional spatial representation of letter domain, using two spatial dimensions, orientation, size and gray-scale shading.

the domain in terms of their derived values on pairs of psychologically meaningful latent variables. A concrete example of this possibility is provided by the depiction of a four-dimensional similarity structure [87] for a set of 16 phonemes, based on their patterns of auditory confusion, which is presented in terms of meaningful stimulus dimensions.

4.7 Arabic Numerals

The seventh case study involves measures of the ‘abstract conceptual similarity’ of the ten Arabic numerals, ‘0’...‘9’, as judged by human subjects, having been pooled across three conditions of stimulus presentation [143].

Featural Visualisation

It is commonly asserted [83, 144] that this type of measure of pairwise similarity is better modelled by the featural than the spatial approach. Accordingly, the results of a




Figure 18: Visualisation of five-dimensional spatial representation of letter domain, showing all ten planar projections.

previous application of an additive clustering technique to this similarity structure were employed [83]. These took the form of a set of 8 derived clusters or features, each with an associated saliency weight, which, together with an additive similarity constant, explained 90.9% of the variance of the data.

Neither the iris domain nor the zoo domain, previously depicted using the ray glyph approach, could claim the principled featural basis provided by this analysis of the numeral domain. The glyph visualisation of the iris domain was based simply on raw measurements, while the zoo domain's featural representation was both subjectively pre-abstracted, and not augmented with featural weightings. Accordingly, Figure 19 presents a ray glyph display based on the featural properties and saliencies for the ten numeral stimuli.

Once again, it is difficult to comprehend the abstract conceptual relationship between the numerals which should be conveyed by this visualisation. With some effort, it might be possible to notice that the numerals 2, 4 and 8 share an important feature, that the numeral 0 is somewhat different from the other numerals, that the numerals 1 and 8 are quite different from each other, and so on. However, the extraction of this information is an effortful process which requires an often myopic local focus, and the global structure of the domain does not seem to be conveyed by the visualisation. Certainly, no sense of the natural cardinal ordering of the stimuli is conveyed by Figure 19.

An established alternative means of displaying a featural representational model [78, 82] is to overlay the derived features upon a two-dimensional spatial presentation of the domain. Figure 20 attempts to do precisely this, using the best fitting two-dimensional




Figure 19: Ray glyph visualisation of additive clustering similarity structure for the Arabic numerals domain.

solution, explaining 60.1% of the variance of the data, as obtained from a multidimensional scaling analysis. The featural clusters are depicted as closed convex shapes which encompass the representative points of the stimuli they contain. It is relatively easy to gain an understanding of the conceptual structure of the numerals using this display by focussing upon each of the clusters, and the underlying spatial configuration conveys aspects, such as cardinality, of the global structure of the domain. While Figure 20 does not provide an indication of the saliencies of each cluster, this has previously been accomplished simply by labelling each boundary [82], and could more naturally be achieved by using weighted lines.

A serious difficulty with this approach to presenting featural representations, however, resides in the strong possibility of having to resort to concave encompassing shapes to impose certain cluster structures upon given spatial configurations. Indeed, this is true of Figure 20, which does not display one of the derived additive clusters, containing the numerals 3, 4, 5, 6 and 7. It is simply not possible to draw a convex encompassing boundary for these stimuli given the fixed spatial representation, and there is no obvious way to modify the underlying configuration to overcome this difficulty. Unfortunately, there are also significant psychological barriers to the introduction of concave bounds, related to their perceptual comprehension in terms of the action of dynamic processes [102]. The use of different encompassing shapes in Figure 20 might be taken as introducing some unintended emphasis on different clusters, and this effect seems likely to magnify by orders




Figure 20: Incomplete visualisation of additive clustering similarity structure of the Arabic numerals using ellipses.

of magnitude if the convexity constraint is relaxed. Ideally, of course, it should be possible to use these perceptual and cognitive biases to communicate the relative weightings of the various clusters, but the current immaturity of understanding in this area prevents the development of principled techniques. In the meantime, therefore, it may be prudent to avoid the depiction of featural cognitive representations using the encompassment scheme shown in Figure 20.

As a workable alternative, Figure 21 presents the derived featural representation of the numeral domain using a novel visualisation approach called 'prototype webs'. In this presentation, as in minimal spanning trees, each cluster is indicated by a 'web' which connects each of the representative points of the stimuli it encompasses. These connecting lines meet at a 'Steiner' point located at the spatial centre of the cluster, which may be regarded as the representational location of the conceptual 'prototype', as in the prototype models of graded conceptual structure discussed earlier. In addition, the relative widths of the connecting lines for each cluster corresponds to the derived saliency of the associated cognitive feature.

The ease with which this visualisation may be comprehended is probably comparable to




Figure 21: Visualisation of additive clustering similarity structure of the Arabic numerals using 'prototype webs'.

that shown in Figure 20, in that aspects of the conceptual structure of the domain may be appreciated through the focussed examination of the individual webs, and the advantages of employing an underlying spatial representation are retained. Given this approximate equivalency, the prototype web approach has the advantage of allowing perceptually similar graphical structures to indicate essentially arbitrary cluster structures.

4.8 Drug Use

The eighth case study considers correlational measures of psychoactive drugs obtained from data which indicated the frequency, on a five point scale, of the use of 13 drugs across 1634 students [145]. Each pairwise correlation between the various drugs was simply treated as constituting a measure of the similarity of those drugs, in the context of frequency of use.

Featural Visualisation

An additive clustering analysis of this similarity structure employing an 'extraction and regularisation' approach [80] generated seven clusters which, when augmented by an additive constant, accounted for 92.6% of the variance of the data. Figure 22 presents a




Figure 22: Visualisation of additive clustering similarity structure of the drug use domain using 'prototype webs'.

prototype web visualisation of the resultant featural representation of the domain¹². Once again, the underlying spatial representation provides some guidance regarding the overall patterns of drug use, and detailed examination of the featural clusters gives an indication of the conceptual associations between specific subsets of the drugs.

Transformational Visualisation

The focussed context of drug usage from which the similarity data for this domain are generated suggests that a transformational representation may be revealing. Since each cluster tends to associate a group of drugs which have the same frequency of use, the pattern of change of drug use should be implicit within the featural representation. By adopting the same transformational approach as was employed for the zoo domain, so that drugs which differ by only one featural cluster are connected, it seems possible that the 'trajectories' by which drug use behaviour changes may be identified.

Figure 23 provides a visualisation of this transformational representation, and seems to indicate not only a separation between legal and illegal drugs, but also a pattern of migration from 'softer' illegal drugs, such as marijuana and hashish, to 'harder' illegal drugs, such as cocaine and heroin. This is a particularly effective and useful means of representing and presenting this domain, and suggests that the development of more sophisticated

¹²Note that the drug labelled 'medication' refers to the abuse of medically prescribed drugs.




Figure 23: Visualisation of transformations for the drug use similarity structure.

techniques for developing transformational representations of similarity structures is warranted.

4.9 Nations

The final case study examines a set of similarity ratings for 17 nations, collected from 557 adults in the mid-to-late 1960s [146]. Based on a series of geographic, demographic, economic and political criteria, the nations considered were Argentina, Brazil, Congo, Cuba, Egypt, France, India, Indonesia, Japan, Nigeria, Phillipines, Poland, Red China, Russia, Spain, United States of America, and Yugoslavia. For each of these nations, subjects were required to judge which 3 of the remaining 16 in the list were the most similar to that nation. The relative frequency, across subjects, with which one nation was identified with another provides a measure of the similarity between those nations. It should be clear, however, that this method of construction allows the possibility of assymetric similarity measures. For example, while 'Indonesia' might prompt geographically motivated responses such as 'Japan' and 'Red China', the prompt 'Japan' might result economic responses such as the 'United States of America', whereas 'Red China' might prompt politically motivated responses such as 'Cuba'.




Figure 24: Asymmetric spatial visualisation of the nations similarity structure. Each nation is labelled by its first four letters.

Assymmetric Spatial Visualisation

Rather than removing this assymetry by, as discussed earlier, considering only a transpose-averaged similarity matrix, Figure 24 attempts to incorporate the patterns of assymetry in a spatial display. Figure 24 augments a presentation of the best-fitting two dimensional spatial representation, determined in relation to the symmetric matrix and explaining only 66.1% of the variance in this data, with a series of 'to' and 'from' arrows and stems indicating deviations from symmetry. The relative lengths of these arrows and stems are proportional to the magnitude of the deviation from symmetry, as measured by the difference between transpose elements in the similarity matrix. For example, the large arrow originating from 'Cuba' and terminating at 'Brazil' indicates that the prompt 'Cuba' resulted in 'Brazil' being listed considerably more often than 'Cuba' itself was listed when the prompt given was 'Brazil'.

Figure 24 appears to maintain the advantages of spatial displays, in that 'communist', 'third-world' and other clusters are immediately perceived, and the incorporation of arrows and stems seems to convey additional information. For example, the status of 'Russia' as a super-power (rather than, say 'Yugoslavia' or 'Poland') is evident from its indicated desire to be more similar to the 'United States' or 'France'. Similarly, the tension between rating 'Cuba' using political and geographical contexts is evident by considering both its location amongst communist nations, and its assymmetric relationship to other South American countries such as 'Argentina' and 'Brazil'.

Despite the obvious potential of presenting assymetric similarity structures, however, the depiction technique shown in Figure 24 constitutes a first attempt at a general method. A range of other visual approaches to conveying assymetry, such as representative geometric shapes with matching intrusions and extrusions, also seem worthy of investigation, and their comparison would be a worthwhile topic for further research.

4.10 Summary

Before commencing a general discussion, it is probably worth summarising some important points arising from the series of case studies just presented. First, the wide variety of means by which measures of pairwise similarity were generated in the case studies should be emphasised. Lists of properties were converted into appropriate measures in several of the cases studies, using Euclidean (iris) and other (zoo) distance metrics, measures based on angular difference (documents), and correlational measures (drug use). Indications of pairwise similarity were also generated from several measures of association, including counts of commonality (senators), raw confusion probabilities (Morse code and English letters), and pooled ratings of similarity (numerals and nations) collected by various means. While there remain a number of well established techniques which have not been covered, the more general point that measures of pairwise similarity are readily generated from a wide variety of raw data seems well made. In particular, both the possibility of measuring psychological similarity using artificial semantic extraction techniques such as n -gram analysis, and using human judgments such as ratings, have been demonstrated by the case studies.

Secondly, some observations should be made regarding the amount of data involved in the case studies, and the implications for the scalability of the visualisation techniques employed. Two of the case studies – the iris measurement and document semantic studies – involved the depiction of something in the order of 100 stimuli, and the ability of the spatial approach to handle significantly increases in the number of stimuli has already been noted. More subtly, however, it needs to be appreciated that many of the case studies manage to consider significant volumes of data, but require the depiction of relatively few stimuli to convey this body of information. For example the transformational visualisation of the drug domain presented in Figure 23 is based on 1634 evaluations of frequency of use for 13 drugs, a total of over 20,000 raw data points. The correlational approach by which the analysis proceeded provides exactly the sort of essence-extracting parsimony that should be achieved by data modelling, and facilitates a final representational structure which is amenable to a simple visualisation. A similar argument could be mounted with respect to the document semantics case study, which is based on more than a megabyte of ‘raw’ ASCII text. The key to the depiction of the clustered structure of the documents shown in Figure 10 is the generation of a similarity structure which accommodates the development of a simple representational model.

5 Future Prospects

When a multidimensional scaling analysis of a given similarity structure indicates that a two-dimensional representation is appropriate, as in the iris, senator, Morse code and




Figure 25: Percentage of variance explained, as a function of the dimensionality of the representational space, for the letter domain.

document case studies, the generation of visualisations of the type shown in Figure 7, 8, 9 and 10 is natural, obvious, and straightforward. This is because the fundamental requirement for effortless and accurate data visualisation – that of canonically aligning a perceptual presentation with an underlying cognitive representation – has been satisfied.

When a two-dimensional spatial representation is inadequate, achieving the necessary match between presentation and representation is more difficult. In fact, there are two possibilities under which it may not be possible to accommodate a similarity structure in two spatial dimensions. It is worth examining both of these possibilities in turn.

5.1 High-Dimensional Spatial Representations

Some similarity structures, such as those involved in the zoo and letter domains, need to be modelled by spatial representations containing more than two dimensions. Figure 25 demonstrates this in relation to the letter stimuli by showing the percentage of variance explained by the best fitting multidimensional configuration as a function of assumed dimensionality of the representational space. In cases such as this, the underlying cognitive representational structure is homogenous – that is, all of the representational dimensions are of the same type – but the usual requirement of depiction in an inherently two-dimensional medium introduces difficulties.

Various established means of visualisation these high-dimensional representations ex-

amined in the case studies, such as the use of glyphs, the introduction of additional perceptual features, or the construction of exhaustive planar projections, do not appear to be particularly effective. The incorporation of size (Figure 13) and gray-scale shading (Figure 15) to convey the third spatial dimension of the zoo representation proved to be moderately effective, but their use appears capable of being subsumed within the special cases of visualising three-dimensional spatial representations using perspective displays (Figure 12). The use of orientation (Figure 14) for three-dimensional structures, or of any combination of perceptual features for higher-dimensional structures (Figure 17) seems to result in displays which defy comprehension. Other than three-dimensional representations, the only other potentially effective visualisation technique found was that of exhaustive planar projections for separable domains.

One possibility worth exploring involves constructing spatial representations which combine subspaces with different metric structures. The use of perceptual features such as size and shading to convey depth is made quantitatively precise through reference to the surface of the two-dimensional medium on which the display is produced and hence, by implication, through reference to the location of the eyes of the observer. Incorporating size or shading in a perspective display, therefore, requires some calculations to gauge the ‘distance’ from the projected three-dimensional representational point to its two-dimensional presentational location, but this distance is independent of the metric structure of the representational space. In other words, under this approach, the presentational size or shading value ascribed to a display is not naturally aligned with the representational value it is intended to convey.

Presumably, one of the reasons why displays such as Figure 13 and Figure 15 are effective is because, serendipitously, they do align presentational and representational values. Both of these visualisations may be regarded as ‘flat’ perspective projections, in which a three-dimensional structure is viewed from directly above. This means that the third dimension, perceptually encoded by shape or shading, is orthogonal to the presentational medium and is, therefore, effectively a uni-dimensional augmentation of the remaining two Euclidean spatial dimensions. This uni-dimensionality, in turn, implies that the metric structure associated with the additional representational dimension is irrelevant, since all Minkowskian distances are the same in one dimension. Therefore, the perceptual relationship between the size or shading used to convey one dimension and the spatial separation used to convey the other two does not matter, since the metric structure of the underlying representation is equally accommodating of all possibilities.

Once higher-dimensional spaces are considered, however, the perceptual relationship between display dimensions should impact upon the metric structure of a derived representation. There seems little doubt that the two perceptual dimensions of spatial separation are integral, and therefore require an underlying Euclidean representation. However, perceptual features such as size, shape, colour, texture, shading and so on are almost certainly separable with respect to spatial separation, but may bear separable, configurational, integral, or other relationships to each other. A broadly reliable determination of these relationships, although not trivial, seems possible, and would provide a series of strong representational constraints from which new multidimensional scaling techniques for data visualisation could be developed. The output of these techniques would be high-dimensional spatial representations in which the metric relationships between the various dimensions were aligned with the perceptual relationships between the features subse-




Figure 26: *Percentage of variance explained, as a function of the dimensionality of the representational space, for the Arabic numeral domain.*

quently used to convey those dimensions. The pursuit of the constraints and techniques needed to realise this possibility is obviously a priority for future research.

5.2 Integrated Similarity Structures

On the other hand, some similarity structures, which are not appropriately modelled by two-dimensional spatial representations, such as those involved in the Arabic numerals and drug use case studies, arise because the data cannot be sufficiently well accommodated by any number of spatial dimensions. An example is provided by the Arabic numeral domain which Figure 26 suggests can only have about 60% of the variance can be explained by a coordinate space representation of any dimensionality. That is, there appears to be an asymptotic level of variance explained which is, presumably, too low to allow the adoption of a solely spatial representation. Unlike the problem with visualising high-dimensional spatial representations, the difficulty here is not one of aligning a presentation with a representation. The problem here is one of developing an adequate representation in the first place.

The obvious means of addressing this problem is to model the given similarity structure using a combination of representational approaches. There is no fundamental reason why models of cognitive representations of stimulus domain must be restricted to either a spatial, or featural, or any other simple approach. Indeed, one of the central concerns of the

structural representational approach is to insist upon the specification of richly structured models of mental representation. The integration of spatial and featural approaches to cognitive representation would seem to constitute a first step towards providing this sort of structure, and certainly offers the promise of being able to combine the representational strengths of both. Whereas, in terms of Figure 2, spatial representations are particularly well suited to modelling the internal structure of mental concepts, featural representations are necessary to capture the interconnectedness of different concepts across levels of abstraction [48, 144].

In this light, it is probably not a coincidence that some of the most effective visualisations presented in the case studies are those which incorporate more than one representational model of the domain. These include Figures 16 and 23, which overlay a feature-based transformational representation upon a spatial representation, and Figures 20, 21 and 22 which overlay an additive clustering featural representation upon a spatial representation.

What is being proposed, however, significantly extends the integration of representational approaches provided by these visualisations, in the sense that these combinations are currently only enacted at a presentational, rather than representational, level. The featural representation of the Arabic numeral stimuli, for example, is derived by an additive clustering technique which operates completely independently from the multidimensional scaling technique by which the spatial representation is generated. The introduction of an integrative inter-dependence between different representational approaches amounts to a fundamentally different, and considerably more sophisticated, type of representational structure, and seems closer to the schematic memory mechanisms which are probably necessary to model human cognition. Exactly what approach should be taken by techniques for developing these types of representations is difficult to determine, although it is reasonable to speculate that there needs to be close interaction between the various representational subcomponents of a schematic structure as they seek to model the similarity relationships existing within a domain.

The possibility of developing integrated representational structures also raises issues related to the correspondence between the perceptual apprehension of similarity in visual displays, and the underlying similarity model employed in a representation. The spatial and featural approaches, for example, differ not only in the types of representations they typically employ – continuous for spatial, and discrete (often binary) for featural – but also in the way in which pairwise measures of similarity are generated from these relationships. Formally, these similarity models are given by Equation 7 for the spatial approach, Equation 8 for the general featural approach and Equation 10 for its additive clustering specialisation. It is evident that these models are fundamentally different, in that the spatial model employs a global approach, in which the modification of one representation affects the similarity of the corresponding stimulus to every other, whereas the featural model is more local, in the sense that ascribing a feature to a stimulus affects its relationship only to the stimuli already having that feature.

The implication of this distinction is that the way in which similarity between presentational components is perceived in a display must be understood, and aligned with the underlying similarity models. For example, the local properties of the featural similarity model seem to correspond to what could be termed ‘associational’ perceptual similarity,

which might be achieved by using sufficiently quantized versions of display features such as size, shape, colour, and so on. It also seems possible that some presentational features might ‘compete’ for attention, which could perhaps be addressed through the introduction of semi-metric spatial representations using Minkowskian r parameters less than 1, as discussed earlier. Clearly, however, there is the need for the development of an empirically consolidated theoretical understanding of these perceptual relationships in the context of data visualisation, before integrated representational structures of the type being proposed could be presented effectively. The pursuit of such an understanding constitutes yet another avenue for future research.

5.3 Conclusion

Both the possibility of veridically displaying high-dimensional spatial representations, and the possibility of generating and displaying integrated similarity structures present major theoretical challenges. The derivation of spatial representations with subspaces operating under different metrics may not be straightforward, and measures of the complexity of such representations also promise to be difficult to determine. Furthermore, even if both of these hurdles are overcome, extensive empirical investigation would be required to align presentational features with representational dimensions. Meanwhile, the derivation of integrated similarity structures requires the development of entirely new approaches to similarity structure modelling. As with any new theoretical enterprise, it is difficult to predict the rate at which progress might be made.

It remains true, however, that any progress on either of these two fronts would significantly enhance the ability to present many domains of interest. While the various representational modelling and visualisation techniques examined in this report often provide a reasonable capability for comprehending disparate bodies of raw information, they all are often subject to clear and fundamental shortcomings. For this reason, partial solutions to the problem of increasing the effectiveness of perceptual displays, and the problem of increasing the sophistication of underlying representations would, no doubt, improve the capabilities of a variety of data visualisation techniques.

Given the goal of data visualisation advocated in this report – that of creating a representational and presentational linkage between artificial information systems and human analysts – some degree of shortcoming in any implemented technique is probably inevitable. What is being sought is effectively a comprehensive understanding of all of human perception and cognition. Whether or not such an understanding is ultimately attainable, approaching data visualisation from the perspective of cognitive and perceptual modelling does ensure that useful theoretical and practical questions are addressed. In this regard, the justification and challenge of psychological approaches to data visualisation is perhaps best summarised by Kosslyn: “principles of perception and memory are a two-edged sword. If they are ignored, a display can be uninterpretable; if respected, it can be read at a glance” [13].

References

1. R. N. Shepard, "Form, formation, and transformation of internal representations," in *Information processing and cognition: The Loyola symposium* (R. L. Solso, ed.), pp. 87–122, Potomac, MD: Erlbaum, 1975.
2. R. N. Shepard, "Psychophysical complementarity," in *Perceptual organization* (M. Kubovy and J. R. Pomerantz, eds.), pp. 279–341, Hillsdale, NJ: Erlbaum, 1981.
3. R. N. Shepard, "Ecological constraints on internal representation: Resonant kinematics of perceiving, imagining, thinking, and dreaming," *Psychological Review*, vol. 91, no. 4, pp. 417–447, 1984.
4. R. N. Shepard, "Evolution of a mesh between principles of the mind and regularities of the world," in *The latest on the best: Essays on evolution and optimality* (Dupré, ed.), pp. 251–275, Cambridge, MA: MIT Press, 1987.
5. D. L. Medin and A. Ortony, "Psychological essentialism," in *Similarity and analogical reasoning* (S. Vosniadou and A. Ortony, eds.), pp. 179–195, New York, NY: Cambridge University Press, 1989.
6. E. W. Averill, "Hidden kind classification," *Psychology of Learning and Motivation*, vol. 29, pp. 437–467, 1993.
7. J. P. Lee and G. S. Grinstein, "Describing visual interactions to the database: Closing the loop between user and data," *SPIE*, vol. 2656, pp. 93–103, 1996.
8. S. Lewandowsky and J. T. Behren, "Visual detection of clusters in statistical maps," in *Proceedings of the American Statistical Society Annual Meeting, Orlando, FL*, 1995.
9. R. M. Pickett, "Perceiving visual texture: A literature survey," Tech. Rep. AMRL-TR-68-12, Behavioral Sciences Laboratory, Aerospace Medical Research Laboratories, 1968.
10. H. Senay and E. Ignatius, "Rules and principles of scientific data visualization," tech. rep., Department of Electrical Engineering and Computer Science, The George Washington University, 1997. [<http://www.cs.gsu.edu/classes/percept/visrules.htm>].
11. B. Schneiderman, *Designing the user interface: Strategies for effective human-computer interaction*. Reading, MA: Addison Wesley, third ed., 1998.
12. I. Lokuge, S. A. Gilbert, and W. Richards, "Structuring information with mental models: A tour of Boston," in *Proceedings of the ACM/SIGCHI Conference on Human Factors in Computing Systems, Vancouver, Canada* (R. Bilger, S. Guest, and M. J. Tauber, eds.), 1996.
13. S. M. Kosslyn, *Elements of Graph Design*. New York, NY: W. H. Freeman, 1994.
14. D. R. Hofstadter, "Foreward," in *Sparse Distributed Memory* (P. Kanerva, ed.), pp. xi–xviii, Cambridge, MA: MIT Press, 1988.

15. L. K. Komatsu, "Recent views of conceptual structure," *Psychological Bulletin*, vol. 112, no. 3, pp. 500-526, 1992.
16. E. Rosch, "Principles of categorization," in *Cognition and categorization* (E. Rosch and B. B. Lloyd, eds.), pp. 27-77, Hillsdale, NJ: Erlbaum, 1978.
17. P. G. Schyns and L. Rodet, "Concept learning," in *Handbook of brain theory and neural networks* (M. A. Arbib, ed.), pp. 234-236, Cambridge, MA: MIT Press, 1995.
18. E. E. Smith, "Concepts and induction," in *Foundations of cognitive science* (M. I. Posner, ed.), pp. 501-526, Cambridge, MA: MIT Press, 1989.
19. J. R. Anderson and L. J. Schooler, "Reflections of the environment in memory," *Psychological Science*, vol. 2, no. 6, pp. 396-408, 1991.
20. P. M. Churchland, *A neurocomputational perspective*. Cambridge, MA: MIT Press, 1989.
21. A. Clark, "Philosophical issues in brain theory and connectionism," in *Handbook of brain theory and neural networks* (M. A. Arbib, ed.), pp. 738-741, Cambridge, MA: MIT Press, 1995.
22. A. Cussins, "The connectionist construction of concepts," in *The philosophy of artificial intelligence* (M. A. Boden, ed.), pp. 368-441, Oxford: Oxford University Press, 1990.
23. J. Fodor, *Psychosemantics: The problem of meaning in the philosophy of mind*. Cambridge, MA: MIT Press, 1987.
24. J. Fodor and Z. Pylyshyn, "Connectionism and cognitive architecture: A critical analysis," *Cognition*, vol. 28, pp. 3-71, 1988.
25. J. S. Bruner, J. J. Goodnow, and G. A. Austin, *A study of thinking*. New York, NY: Wiley, 1956.
26. R. Wille, "Concept lattices and conceptual knowledge systems," *Computers in mathematics with applications*, vol. 23, no. 6-9, pp. 493-515, 1992.
27. E. Rosch and C. B. Mervis, "Family resemblances: Studies in the internal structure of categories," *Cognitive Psychology*, vol. 7, pp. 573-605, 1975.
28. D. L. Medin and M. M. Schaffer, "Context theory of classification," *Psychological Review*, vol. 85, pp. 207-238, 1978.
29. R. M. Nosofsky, "Choice, similarity, and the context theory of classification," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, vol. 10, no. 1, pp. 104-114, 1984.
30. D. L. Medin, "Concepts and conceptual structure," *American Psychologist*, vol. 44, no. 12, pp. 1469-1481, 1989.
31. E. Rosch, C. B. Mervis, W. D. Gray, D. M. Johnson, and P. Boyes-Braem, "Basic objects in natural categories," *Cognitive Psychology*, vol. 8, pp. 382-439, 1976.

32. M. A. Gluck and G. H. Bower, "From conditioning to category learning: An adaptive network model," *Journal of Experimental Psychology: General*, vol. 117, pp. 227-247, 1988.
33. G. B. Booch, *Object-oriented analysis and design*. Santa Clara, CA: Benjamin Cummings, second ed., 1994.
34. R. N. Shepard, "George Miller's data and the development of methods for representing cognitive structures," in *The making of cognitive science* (W. Hirst, ed.), pp. 45-70, Cambridge, MA: Cambridge University Press, 1988.
35. R. L. Goldstone, "Similarity," in *MIT encyclopedia of the cognitive sciences* (R. A. Wilson and F. C. Keil, eds.), Cambridge, MA: MIT Press, in press.
36. D. Brodbeck, M. Chalmer, A. Lunzer, and P. Cotture, "Domesticating Bead: Adapting an information visualization system to a financial institution," in *IEEE Symposium on Information Visualization* (J. Dill and N. Gershon, eds.), (Los Alamitos, CA), pp. 73-80, IEEE Computer Society, 1997.
37. T. F. Cox and M. A. A. Cox, *Multidimensional scaling*. London: Chapman and Hall, 1994.
38. R. N. Shepard, "Stimulus and response generalization: A stochastic model relating generalization to distance in psychological space," *Psychometrika*, vol. 22, no. 4, pp. 325-345, 1957.
39. R. N. Shepard, "Psychological representation of speech sounds," in *Human communication: A unified view* (E. E. David and P. B. Denes, eds.), pp. 67-113, New York, NY: McGraw Hill, 1972.
40. R. M. Nosofsky, "Similarity scaling and cognitive process models," *Annual Review of Psychology*, vol. 43, pp. 25-53, 1992.
41. M. D. Lee, "The connectionist construction of psychological spaces," *Connection Science*, vol. 9, no. 4, pp. 323-351, 1997.
42. R. N. Shepard, "The analysis of proximities: Multidimensional scaling with an unknown distance function. II.," *Psychometrika*, vol. 27, no. 3, pp. 219-246, 1962.
43. T. F. Cox and M. A. A. Cox, "Multidimensional scaling on a sphere," *Communications in Statistics: Theory and Methods*, vol. 20, no. 9, pp. 2943-2953, 1991.
44. H. Lindman and T. Caelli, "Constant curvature Riemannian scaling," *Journal of Mathematical Psychology*, vol. 17, pp. 89-109, 1978.
45. W. R. Garner, *The processing of information and structure*. Potomac, MD: Erlbaum, 1974.
46. R. N. Shepard, "Integrality versus separability of stimulus dimensions: From an early convergence of evidence to a proposed theoretical basis," in *The perception of structure: Essays in honor of Wendell R Garner* (J. R. Pomerantz and G. L. Lockhead, eds.), pp. 53-71, Washington, DC: American Psychological Association, 1991.

47. R. N. Shepard, "Toward a universal law of generalization for psychological science," *Science*, vol. 237, pp. 1317-1323, 1987.
48. R. N. Shepard, "Perceptual-cognitive universals as reflections of the world," *Psychonomic Bulletin & Review*, vol. 1, no. 1, pp. 2-28, 1994.
49. I. J. Myung and R. N. Shepard, "Maximum entropy inference and stimulus generalization," *Journal of Mathematical Psychology*, vol. 40, pp. 342-347, 1996.
50. J. B. Kruskal, "Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis," *Psychometrika*, vol. 29, no. 1, pp. 1-27, 1964.
51. I. Gati and A. Tversky, "Representations of qualitative and quantitative dimensions," *Journal of Experimental Psychology: Human Perception and Performance*, vol. 8, no. 2, pp. 325-340, 1982.
52. I. Borg and J. Lingoes, *Multidimensional similarity structure analysis*. New York, NY: Springer Verlag, 1987.
53. R. N. Shepard, "Representation of structure in similarity data: Problems and prospects," *Psychometrika*, vol. 39, no. 4, pp. 373-422, 1974.
54. J. D. Cohen, "Drawing graphs to convey proximity: An incremental arrangement method," *ACM Transactions On Computer-Human Interaction*, vol. 4, no. 3, pp. 197-229, 1997.
55. M. L. Davison, *Multidimensional scaling*. New York, NY: Wiley, 1983.
56. J. W. Sammon, "A nonlinear mapping for data structure analysis," *IEEE Transactions on Computers*, vol. 18, pp. 401-409, 1969.
57. S. S. Schiffman, M. L. Reynolds, and F. W. Young, *Introduction to multidimensional scaling*. New York, NY: Academic Press, 1981.
58. R. N. Shepard, "The analysis of proximities: Multidimensional scaling with an unknown distance function. I," *Psychometrika*, vol. 27, no. 2, pp. 125-140, 1962.
59. M. D. Lee, "Determining the dimensionality of multidimensional scaling representations for cognitive modeling." Manuscript submitted for publication, 1998.
60. J. Mao and A. K. Jain, "Artificial neural networks for feature extraction and multivariate data projection," *IEEE Transactions on Neural Networks*, vol. 6, no. 2, pp. 296-317, 1995.
61. P. Demartines and J. Héroult, "Curvilinear component analysis: A self-organizing neural network for nonlinear mapping of data sets," *IEEE Transactions on Neural Networks*, vol. 8, no. 1, pp. 148-154, 1997.
62. D. Lowe and M. E. Tipping, "Feed-forward neural networks and topographic mappings for exploratory data analysis," *Neural Computing and Applications*, vol. 4, pp. 83-95, 1996.

63. A. R. Webb, "Multidimensional scaling by iterative majorization using radial basis functions," *Pattern Recognition*, vol. 28, no. 5, pp. 753-759, 1995.
64. J. A. Anderson, *An introduction to neural networks*. Cambridge, MA: MIT Press, 1995.
65. D. E. Rumelhart, "The architecture of mind: A connectionist approach," in *Foundations of cognitive science* (M. I. Posner, ed.), pp. 133-159, Cambridge, MA: MIT Press, 1989.
66. G. E. Hinton, J. L. McClelland, and D. E. Rumelhart, "Distributed representations," in *Parallel distributed processing: Explorations in the microstructure of cognition, vol. 1* (D. E. Rumelhart and J. L. McClelland, eds.), pp. 77-109, Cambridge, MA: MIT Press, 1986.
67. D. E. Rumelhart, P. Smolensky, J. L. McClelland, and G. E. Hinton, "Schemata and sequential thought processes in PDP models," in *Parallel distributed processing: Explorations in the microstructure of cognition, vol. 2* (J. L. McClelland and D. E. Rumelhart, eds.), pp. 7-57, Cambridge, MA: MIT Press, 1986.
68. S. Kaski, "Data exploration using self-organizing maps," 1997. Acta Polytechnica Scandinavia, Mathematics, Computing and Management in Engineering Series No. 82, Espoo, Finland.
69. S. Kaski, "Computationally efficient approximation of a probabilistic model for document representation in the WEBSOM full-text analysis method," *Neural Processing Letters*, vol. 5, pp. 139-151, 1997.
70. T. Kohonen, "Self-organized formation of topologically correct feature maps," *Biological Cybernetics*, vol. 43, pp. 59-69, 1982.
71. T. Kohonen, "The self-organizing map," *Proceedings of the IEEE*, vol. 78, no. 9, pp. 1464-1480, 1990.
72. R. Durbin and D. Willshaw, "An analogue approach to the travelling salesman problem using an elastic net method," *Nature*, vol. 326, pp. 689-691, 1987.
73. A. L. Yuille, "Constrained optimization and the elastic net," in *Handbook of brain theory and neural networks* (M. A. Arbib, ed.), pp. 250-255, Cambridge, MA: MIT Press, 1995.
74. C. M. Bishop, M. Svensén, and C. K. I. Williams, "GTM: The generative topographic mapping," *Neural Computation*, vol. 10, no. 1, pp. 215-234, 1998.
75. A. Flexer, "Limitations of self-organizing maps for vector quantization and multidimensional scaling," 1996. Unpublished manuscript, Austrian Research Institute for Artificial Intelligence.
76. H. Ritter, "Self-organising maps for internal representations," *Psychological Research*, vol. 52, no. 2-3, pp. 128-136, 1990.
77. A. Tversky, "Features of similarity," *Psychological Review*, vol. 84, no. 4, pp. 327-352, 1977.

78. P. Arabie and J. D. Carroll, "MAPCLUS: A mathematical programming approach to fitting the ADCLUS model," *Psychometrika*, vol. 45, no. 2, pp. 211-235, 1980.
79. A. Chaturvedi and J. D. Carroll, "An alternating combinatorial optimization approach to fitting the INDCLUS and generalized INDCLUS models," *Journal of Classification*, vol. 11, pp. 155-170, 1994.
80. M. D. Lee, "An extraction and regularization approach to additive clustering." Manuscript submitted for publication, 1997.
81. B. G. Mirkin, "Additive clustering and qualitative factor analysis methods for similarity matrices," *Journal of Classification*, vol. 4, pp. 7-31, 1987.
82. R. N. Shepard and P. Arabie, "Additive clustering representations of similarities as combinations of discrete overlapping properties," *Psychological Review*, vol. 86, no. 2, pp. 87-123, 1979.
83. J. B. Tenenbaum, "Learning the structure of similarity," in *Advances in neural information processing systems* (D. S. Touretzky, M. C. Mozer, and M. E. Hasselmo, eds.), vol. 8, pp. 3-9, Cambridge, MA: MIT Press, 1996.
84. M. D. Lee, "On the complexity of additive clustering models." Manuscript submitted for publication, 1997.
85. J. E. Corter, *Tree Models of Similarity and Association*. Thousand Oaks, CA: Sage, 1996.
86. S. Sattath and A. Tversky, "Additive similarity trees," *Psychometrika*, vol. 42, pp. 319-345, 1977.
87. R. N. Shepard, "Multidimensional scaling, tree-fitting, and clustering," *Science*, vol. 210, pp. 390-398, 1980.
88. M. Minsky, "A framework for representing knowledge," in *The psychology of computer vision* (P. H. Winston, ed.), pp. 211-277, New York, NY: McGraw-Hill, 1975.
89. R. C. Schank and R. P. Abelson, *Scripts, plans, goals and understanding*. Hillsdale, NJ: Erlbaum, 1977.
90. D. E. Rumelhart, "Schemata: The building block of cognition," in *Theoretical issues in reading comprehension* (R. Spiro, B. Bruce, and W. Brewer, eds.), pp. 33-58, Hillsdale, NJ: Erlbaum, 1980.
91. D. R. Hofstadter, *Gödel, Escher, Bach: An eternal golden braid*. New York, NY: Basic Books, 1979.
92. D. R. Hofstadter, *Metamagical themas: Questing for the essence of mind and pattern*. New York, NY: Basic Books, 1985.
93. R. L. Goldstone, "Similarity, interactive activation, and mapping," *Journal of Experimental Psychology: Learning, Memory and Cognition*, vol. 20, pp. 3-28, 1994.

94. M. D. Lee, "Neural feature abstraction from judgments of similarity," *Neural Computation*, in press.
95. P. Smolensky, "On the proper treatment of connectionism – again," *Behavioral and Brain Sciences*, vol. 11, pp. 59–74, 1988.
96. E. H. Carlton and R. N. Shepard, "Psychologically simple motions as geodesic paths I. Asymmetric objects," *Journal of Mathematical Psychology*, vol. 34, pp. 127–188, 1990.
97. D. Vickers, "An Erlanger programme for psychology." Paper presented at Symposium on Chaos Theory Applied to Social Science Research, Fourth International Social Science Methodology Conference, Colchester, Essex, U.K., 1996.
98. D. Vickers and M. D. Lee, "Toward a dynamic connectionist model of memory," *Behavioral and Brain Sciences*, vol. 20, no. 1, pp. 40–41, 1997.
99. R. N. Shepard and L. A. Cooper, *Mental images and their transformations*. Cambridge, MA: MIT Press, 1982.
100. E. H. Carlton and R. N. Shepard, "Psychologically simple motions as geodesic paths II. Symmetric objects," *Journal of Mathematical Psychology*, vol. 34, pp. 189–228, 1990.
101. J. J. Freyd, "Dynamic mental representations," *Psychological Review*, vol. 94, no. 4, pp. 427–438, 1987.
102. M. Leyton, *Symmetry, causality, mind*. Cambridge, MA: MIT Press, 1992.
103. S. E. Palmer, "Goodness, Gestalt, groups, and Garner," in *The perception of structure: Essays in honor of Wendell R Garner* (J. R. Pomerantz and G. L. Lockhead, eds.), pp. 23–39, Washington, DC: American Psychological Association, 1991.
104. W. Richards and J. J. Koenderink, "Trajectory mapping: A new nonmetric scaling technique," *Perception*, vol. 24, pp. 1315–1331, 1995.
105. R. A. Fisher, "The use of multiple measurements in taxonomic problems," *Annals of Eugenics*, vol. 7, pp. 179–184, 1936.
106. W. Ribarsky, E. Ayers, J. Eble, and S. Mukherjea, "Glyphmaker: Creating customized visualizations of complex data," *IEEE Computer*, vol. 27, no. 7, pp. 57–64, 1997.
107. H. Chernoff, "The use of faces to represent points in k-dimensional space graphically," *Journal of the American Statistical Association*, vol. 68, pp. 361–368, 1973.
108. H. Chernoff and M. H. Rizvi, "Effect on classification error of random permutations of features in representing multivariate data by faces," *Journal of American Statistical Association*, vol. 70, pp. 548–554, 1975.
109. J. M. Chalmers, W. S. Cleveland, B. Kleiner, and P. A. Tukey, *Graphical Methods for Data Analysis*. Boston, MA: Duxbury Press, 1983.

110. E. R. Tufte, *The Visual Display of Quantitative Information*. Cheshire, CT: Graphics Press, 1983.
111. R. N. Shepard, "The perceptual organization of colors: An adaptation to regularities of the terrestrial world?," in *The Adapted Mind: Evolutionary Psychology and the Generation of Culture* (J. H. Barkow, L. Cosmides, and J. Tooby, eds.), pp. 211-236, New York, NY: Oxford University Press, 1992.
112. D. Travis, *Effective Color Displays: Theory and Practice*. New York, NY: Academic Press, 1991.
113. D. Vickers, *Decision Processes in Visual Perception*. New York, NY: Academic Press, 1979.
114. T. E. Goldsmith and R. W. Schvaneveldt, "Facilitating multiple-cue judgments with integral information displays," in *Human Factors in Computer Systems* (J. Thomas and M. Schneider, eds.), Norwood, NJ: Ablex, 1985.
115. S. S. Stevens, *Psychophysics*. New York, NY: Wiley, 1975.
116. J. Mackinlay, "Automatic design of graphical presentations," Tech. Rep. Stan-CS-86-1038, Department of Computer Science, Stanford University, 1986.
117. J. C. Baird, *Psychophysical Analysis of Visual Space*. New York, NY: Permagon Press, 1970.
118. C. H. Yu, "Visualization techniques of different dimensions," 1998. [<http://seamon-key.ed.asu.edu/behrens/asu/reports/compre/comp1.html#RTFToC11>].
119. J. Bertin, *Semiology of Graphics: Diagrams, Networks, Maps*. Madison, WI: University of Wisconsin Press, 1983. Translated by W. J. Berg.
120. H. C. Romesburg, *Cluster analysis for researchers*. Belmont, CA: Lifetime Learning Publications, 1984.
121. E. Z. Rothkopf, "A measure of stimulus similarity and errors in some paired-associate learning tasks," *Journal of Experimental Psychology*, vol. 53, pp. 94-101, 1957.
122. M. D. Lee, "Consequential region scaling," 1997. Manuscript submitted for publication.
123. D. S. Touretzky, M. C. Mozer, and M. E. Hasselmo, eds., *Neural Information Processing Systems*, (Cambridge, MA), MIT Press, 1996.
124. M. Damashek, "Gauging similarity with n-grams: Language-independent categorization of text," *Science*, vol. 267, pp. 843-848, 1995.
125. S. Lewandowsky and I. Spence, "Discriminating strata in scatterplots," *Journal of the American Statistical Association*, vol. 84, pp. 682-688, 1989.
126. G. E. Legge, Y. Gu, and A. Luebker, "Efficiency of graphical perception," *Perception & Psychophysics*, vol. 46, pp. 365-374, 1989.

127. C. J. Merz and P. M. Murphy, "UCI repository of machine learning databases," 1996. [<http://www.ics.uci.edu/mlearn/MLRepository.html>].
128. S. Pinker, "A theory of graph comprehension," in *Artificial Intelligence and the Future of Testing* (R. Freedle, ed.), pp. 73–126, Hillsdale, NJ: Erlbaum, 1990.
129. D. MacGregor and P. Slovic, "Graphic representation of judgmental information," *Human-Computer Interaction*, vol. 2, pp. 179–200, 1986.
130. D. Simkin and R. Hastie, "An information processing analysis of graph perception," *Journal of the American Statistical Association*, vol. 82, pp. 454–465, 1987.
131. F. Attneave, "Some informational aspects of visual perception," *Psychological Review*, vol. 61, pp. 183–193, 1954.
132. N. Naikar, "Human factors considerations in perspective display design," Tech. Rep. in press, DSTO, 1998.
133. C. D. Wickens, S. Todd, and K. Seidler, "Three-dimensional displays: Perception, implementation, applications," Tech. Rep. 89-01., CSERIAC State-of-the-Art, 1989.
134. M. C. Escher, *M. C. Escher: The Graphic Work*. Berlin, Germany: Benedikt Taschen, 1990.
135. W. S. Cleveland and R. McGill, "Graphical perception: The visual decoding of quantitative information on graphical displays of data," *Journal of the Royal Statistical Society: Series A*, vol. 150, no. 3, pp. 192–229, 1987.
136. L. V. Peterson and W. Schramm, "How accurately are different kinds of graphs read?," *Audio-Visual Communication Review*, vol. 2, pp. 178–189, 1954.
137. I. Spence and S. Lewandowsky, "Displaying proportions and percentages," *Applied Cognitive Psychology*, vol. 5, pp. 61–77, 1991.
138. D. J. Cohen, "Visual detection and perceptual independence: Assessing color and form," *Perception & Psychophysics*, vol. 59, no. 4, pp. 623–635, 1997.
139. M. K. Coleman and D. S. Parker, "Aesthetics-based graph layout for human consumption," *Software-Practice and Experience*, vol. 26, no. 12, pp. 1415–1438., 1996.
140. W. Köhler, *Gestalt Psychology*. New York, NY: Liveright, 1929.
141. G. J. Goodhill, M. W. Simmen, and D. J. Willshaw, "An evaluation of the use of multidimensional scaling for understanding brain connectivity," Tech. Rep. EUCCS/RP-63, Centre for Cognitive Science, University of Edinburgh, 1994.
142. E. J. Gibson, H. Osser, W. Schiff, and J. Smith, "An analysis of critical features of letters, tested by a confusion matrix," Tech. Rep. Cooperative Research Project No. 639, U.S. Office of Education, 1963.
143. R. N. Shepard, D. W. Kilpatrick, and J. P. Cunningham, "The internal representation of numbers," *Cognitive Psychology*, vol. 7, pp. 82–138, 1975.

144. A. Tversky and J. W. Hutchinson, "Nearest neighbor analysis of psychological spaces," *Psychological Review*, vol. 93, no. 1, pp. 3-22, 1986.
145. G. L. Huba, J. A. Wingard, and P. M. Bentler, "A comparison of two latent variable causal models for adolescent drug use," *Journal of Personality and Social Psychology*, vol. 40, no. 1, pp. 180-193, 1981.
146. J. P. Robinson and R. Hefner, "Multidimensional differences in public and academic perceptions of nations," *Journal of Personality and Social Psychology*, vol. 7, no. 3, pp. 251-259, 1967.

THIS PAGE IS INTENTIONALLY BLANK

DISTRIBUTION LIST

Psychological Approaches To Data Visualisation

Michael D. Lee and
Douglas Vickers

Number of Copies

DEFENCE ORGANISATION

S&T Program

Chief Defence Scientist	}	
FAS Science Policy	}	1
AS Science Corporate Management	}	
Director General Science Policy Development		1
Counsellor, Defence Science, London		Doc Data Sht
Counsellor, Defence Science, Washington		Doc Data Sht
Scientific Adviser to MRDC, Thailand		Doc Data Sht
Director General Scientific Advisers and Trials	}	
Scientific Adviser Policy and Command	}	1
Navy Scientific Adviser		Doc Data Sht
Scientific Adviser, Army		Doc Data Sht
Air Force Scientific Adviser		1
Director Trials		1

Aeronautical and Maritime Research Laboratory

Director, Aeronautical and Maritime Research Laboratory		1
Research Leader, Simulation & Human Factors		1
Dr. Gavan Lintern, AOD		1
Dr. Darryl McMahon, MOD		1

Electronics and Surveillance Research Laboratory

Director, Electronics and Surveillance Research Laboratory		1
Chief, Communications Divison		1
Research Leader, Secure Communications		1
Research Leader, Military Information Networks		1
Head, Intelligent Networks, CD		1
Head, C3I Networks, CD		1
Head, Human Systems Integration, ITD		1
Head, Command and Control Australian Theatre, ITD		1
Mr. Victor Demczuk, LOD		1
Dr. Michael Webb, CD		1
Dr. Michael Lee (Author), CD		5

DSTO Libraries

Library Fishermens Bend	1
Library Maribyrnong	1
Library Salisbury	2
Australian Archives	1
Library, MOD, Pyrmont	Doc Data Sht

Capability Development Division

Director General Maritime Development	Doc Data Sht
Director General Land Development	Doc Data Sht
Director General Aerospace Development	Doc Data Sht
Director General C3I Development	Doc Data Sht

Army

ABCA Office, G-1-34, Russell Offices, Canberra	4
NAPOC QWG Engineer NBCD c/- DENGERS-A, HQ Engineer Centre Liverpool Military Area, NSW 2174	Doc Data Sht

Intelligence Program

Director General Research & Development, DSD	1
Defence Intelligence Organisation	1

Corporate Support Program (libraries)

Officer in Charge, TRS, Defence Regional Library, Canberra	1
Officer in Charge, Document Exchange Centre	Doc Data Sht
Additional copies for DEC for exchange agreements	
US Defense Technical Information Center	2
UK Defence Research Information Centre	2
Canada Defence Scientific Information Service	1
NZ Defence Information Centre	1
National Library of Australia	1

UNIVERSITIES AND COLLEGES

Australian Defence Force Academy Library	1
Head of Aerospace and Mechanical Engineering, ADFA	1
Deakin University Library, Serials Section (M List)	1
Senior Librarian, Hargrave Library, Monash University	1
Librarian, Flinders University	1
Librarian, Barr-Smith Library, University of Adelaide	1
Dr. Douglas Vickers, Department of Psychology, University of Adelaide (Author)	5

OTHER ORGANISATIONS

NASA (Canberra)	1
Australian Government Publishing Service	1
The State Library of South Australia	1
Parliamentary Library of South Australia	1

ABSTRACTING AND INFORMATION ORGANISATIONS

INSPEC: Acquisitions Section Institution of Electrical Engineers	1
Engineering Societies Library, US	1
Materials Information, Cambridge Science Abstracts, US	1
Documents Librarian, The Center for Research Libraries, US	1

INFORMATION EXCHANGE AGREEMENT PARTNERS

Acquisitions Unit, Science Reference and Information Service, UK	1
Library – Exchange Desk, National Institute of Standards and Technology, US	1

SPARES

DSTO Salisbury Research Library	10
---------------------------------	----

Total number of copies:	74
--------------------------------	-----------

Page classification: UNCLASSIFIED

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA				1. CAVEAT/PRIVACY MARKING	
2. TITLE Psychological Approaches To Data Visualisation			3. SECURITY CLASSIFICATION Document (U) Title (U) Abstract (U)		
4. AUTHOR(S) Michael D. Lee and Douglas Vickers			5. CORPORATE AUTHOR Electronics and Surveillance Research Laboratory PO Box 1500 Salisbury, South Australia, Australia 5108		
6a. DSTO NUMBER DSTO-RR-0135		6b. AR NUMBER AR-010-587		6c. TYPE OF REPORT Research Report	
				7. DOCUMENT DATE July, 1998	
8. FILE NUMBER E8709-7-11		9. TASK NUMBER DST 96/149		10. SPONSOR DSTO	
				11. No OF PAGES 53	
				12. No OF REFS 146	
13. DOWNGRADING / DELIMITING INSTRUCTIONS Not Applicable			14. RELEASE AUTHORITY Chief, Communications Division		
15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT <i>Approved For Public Release</i> OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE CENTRE, DIS NETWORK OFFICE, DEPT OF DEFENCE, CAMPBELL PARK OFFICES, CANBERRA, ACT 2600					
16. DELIBERATE ANNOUNCEMENT No Limitations					
17. CITATION IN OTHER DOCUMENTS No Limitations					
18. DEFTEST DESCRIPTORS Perception (psychology) Visualisation Cognition Comprehension Integration					
19. ABSTRACT The aim of 'data visualisation' is to display a body of information in a way which allows accurate and effortless human comprehension and analysis. Accordingly, the development of data visualisation techniques should be constrained by an understanding of both human perception and cognition. This report develops and examines a psychological framework for the development of data visualisation techniques based on the notion of similarity structure modelling. Through a series of case studies, a range of established approaches to data visualisation is reviewed and evaluated within this framework, and a number of suggestions for the development of new techniques is made.					

Page classification: UNCLASSIFIED